



Sun StorEdge™ T3+ Array Administrator's Manual

Version 2.1 Controller Firmware

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

Part No. 816-4770-10
April 2002, Revision A

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, SunService, StorTools, and, Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a nonexclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

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

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 Etats-Unis. Tous droits réservés.

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

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, SunSolve, SunService, StorTools, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" 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 xv

- 1. Overview of the Sun StorEdge T3+ Array** 1
 - Product Description 2
 - Array Components 4
 - Disk Drive 4
 - Array Controller Card 5
 - Unit Interconnect Cards (UIC) 7
 - Power and Cooling Unit 10
 - Array Features 11
 - Array Architecture 12

- 2. Array Firmware 2.1 New Features Concepts** 15
 - Firmware Version 2.1 New Features List 16
 - Firmware Version 2.1 New Features Concepts 16
 - Volume Slicing, Mapping, and Masking 16
 - Auto Upgrade 20
 - Split Loop Mode 20
 - Fabric Point-to-Point Support 20
 - Persistent Group Reservation 20

Network Time Protocol 21

Auto-Versioning 21

3. Array Administration Overview 23

▼ To Display Commands 24

▼ To Display Command Syntax 24

FRU Identifiers 25

4. Resetting Global Parameters 27

▼ To Set the Cache Block Size 28

▼ To Enable Partner Group Multipathing 30

▼ To Set the Cache Mode 30

▼ to Enable Mirrored Cache 31

▼ To Set the LUN Reconstruction Rate 32

▼ To Perform Volume Verification 33

▼ To Set Cache Read-Ahead Threshold 34

▼ To Set the IP Address 35

5. Reconfiguring a Volume 37

Reconfiguration Limitations 38

▼ To Delete a Logical Volume 38

▼ To Create a Logical Volume 39

▼ To Label a Volume 41

6. Monitoring the Array 45

▼ To Determine Failover 46

▼ To Check the Drive Status 47

▼ To Check the Hot-Spare 48

▼ To Check Data Parity 48

▼ To Check the Battery 49

- ▼ To Display the FRU Information 50
- ▼ to Check the FRU Status 51
- 7. Array Configuration 55**
 - Configuring the Array for Remote System Logging 56
 - ▼ To Transfer Array Files to the Management Host 56
 - ▼ To Edit the Array /etc/syslog.conf File 58
 - ▼ To Edit the Array /etc/hosts File 59
 - ▼ To Transfer Files Back to the Array 60
 - ▼ To Edit the Management Host /etc/syslog.conf File 61
 - Configuring the Array for SNMP Notification 63
 - ▼ To Transfer the Array Files to the Management Host 63
 - ▼ To Edit the Array /etc/syslog.conf File 65
 - ▼ To Edit the Array /etc/hosts File 67
 - ▼ To Transfer Files Back to the Array 68
- 8. Volume Slicing, Mapping, and Masking Configuration 71**
 - Volume Configuration Limitations 72
 - Volume Slicing Limitations 73
 - T3 Array Configuration Considerations 73
 - Specifying Volume Slicing 75
 - ▼ To Enable Volume Slicing, Mapping, and Masking Operations 76
 - ▼ To Reconfigure a Volume by Adding Slices After Volume Slicing Has Been Enabled 79
 - ▼ To Create a Volume Slice 88
 - ▼ To Remove a Volume Slice 88
 - Specifying LUN Mapping 89
 - ▼ To Display the LUN Map 89
 - ▼ To Add an Entry to the LUN Map 90
 - ▼ To Remove an Entry From the LUN Map 91

Specifying LUN Masking 92

- ▼ To Find the Default LUN Permissions and System-Assigned WWNs 93
- ▼ To Set a Specific LUN Permission 94
- ▼ To View All Registered WWNs 94

Setting LUN Permission with Host Multipathing 95

- ▼ To Change the Default LUN Permissions 96
- ▼ To Define a Host WWN Group 97
- ▼ To Set the Permissions of a LUN for All Members of a Group 98
- ▼ To Remove the Registered Status of a WWN 100
- ▼ To Remove a WWN From a Specific Group 101
- ▼ To Remove All WWNs From a Specific Group 102

9. Network Time Protocol Configuration 105

Specifying Network Time Protocol 106

- ▼ To Set Up the NTP Server 106

10. Fibre Channel Topology Configuration 109

Specifying Fibre Channel Topology 110

- ▼ To Determine the Topology Firmware Mode 110
- ▼ To Change the Topology Firmware Mode 111

11. Advanced Array Troubleshooting 113

- ▼ To Identify FRU Failures 114

Replacing FRUs 119

- ▼ To Verify the FRU Firmware Level 119
- ▼ To Install the Replacement FRU 121

Controller Card 123

Power and Cooling Unit 123

Interconnect Card 123

Chassis and Centerplane 124

▼ To Verify the FRU Replacement	124
▼ To Identify Miscabled Partner Groups	126
Identifying Data Channel Failures	130
12. CLI Commands	131
Complete List of CLI Commands	132
FRU Identifiers	135
Code Conventions	135
General Commands	136
General Command Descriptions	137
Volume slicing, Mapping, and Masking Commands	158
Volume Slicing Command Descriptions	158
LUN Mapping Command Descriptions	165
LUN Masking Commands Descriptions	171
Fibre Channel Topology Command	185
Network Time Protocol Commands	189
13. syslog Error Messages	201
Error Message Syntax	202
Message Types	202
FRU Identifiers	203
Error Message Listing	203
Warning Message Listing	208
Reset Log Types	219
Command-Line Error Messages	220
RAID Errors and Other Common Errors	221
Categories of RAID Errors	223
Investigating RAID Errors	225
Port Errors	225

Interconnect Card and Other FRU Errors	226
Other Errors	228
Examples	230
Error Messages	230
FRU Level Errors	230
Warning Messages	231
Cache Parity Error	232
Software Assertion or Panic	232
Unrecoverable Drive Error	233
Recoverable Drive Error	233
Notice	234
Glossary	235

Figures

- FIGURE 1-1 Sun StorEdge T3+ Array 3
- FIGURE 1-2 Disk Drives (Front View) 5
- FIGURE 1-3 Controller Card (Rear View) 6
- FIGURE 1-4 Sun StorEdge T3+ Array Controller Card Front Panel 6
- FIGURE 1-5 Sun StorEdge T3+ Array Controller Card and Ports 7
- FIGURE 1-6 Interconnect Card and Ports 8
- FIGURE 1-7 Interconnect Cards (Rear View) 9
- FIGURE 1-8 Power and Cooling Units (Rear View) 10
- FIGURE 2-1 Two Possible Mapping Configurations for a Single T3+ Array Unit With Controller Firmware Version 2.0.1 17
- FIGURE 2-2 Two Possible Mapping Configurations for a Single T3+ Array Unit with Controller Firmware Version 2.1 17
- FIGURE 2-3 Two Possible Mapping Configurations for a Single T3+ Array Unit With Controller Firmware Version 2.1 and Volume Slicing Enabled 18
- FIGURE 2-4 A Slicing Example for a T3+ Array Enterprise Group With Each T3+ Unit Having Two Volumes, Controller Firmware Version 2.1, and Volume Slicing Enabled 19
- FIGURE 8-1 Sun StorEdge T3+ Array Enterprise Configuration Example 79
- FIGURE 8-2 Host Multipathing Configuration 95
- FIGURE 11-1 MAC Address on Pull-out Tab 127
- FIGURE 11-2 Power Switch Locations 129

Tables

TABLE 1-1	Sun StorEdge T3+ Array Key Features—73-Gbyte Drive	11
TABLE 3-1	FRU Identifiers	25
TABLE 5-1	Volume Configuration Examples	39
TABLE 6-1	Drive Status Codes	47
TABLE 6-2	Possible FRU States and Definitions	52
TABLE 7-1	Message Categories	58
TABLE 7-2	Message Categories	66
TABLE 8-1	T3+ Array Volume and Drive Configurations Available per RAID Level	74
TABLE 10-1	<code>fc_topology</code> mode settings	111
TABLE 12-1	CLI Commands Listed in Alphabetical Order for T3+ Array with Controller Firmware Version 2.1	132
TABLE 12-2	FRU Identifiers	135
TABLE 12-3	Text Code Conventions	135
TABLE 12-4	<code>boot</code> Command Options and Parameters	137
TABLE 12-5	<code>disable</code> Command Options and Parameters	138
TABLE 12-6	<code>disk</code> Command Options and Parameters	138
TABLE 12-7	<code>enable</code> Command Options and Parameters	139
TABLE 12-8	<code>ep</code> Command Options and Parameters	139
TABLE 12-9	<code>fru</code> Command Options and Parameters	141
TABLE 12-10	<code>help</code> Command Options and Parameters	141
TABLE 12-11	<code>id</code> Command Options and Parameters	142

TABLE 12-12	<code>logger</code> Command Options and Parameters	143
TABLE 12-13	<code>lpc</code> Command Options and Parameters	144
TABLE 12-14	<code>ofdg</code> Command Options and Parameters	145
TABLE 12-15	<code>port</code> Command Options and Parameters	146
TABLE 12-16	<code>proc</code> Command Options and Parameters	147
TABLE 12-17	<code>refresh</code> Command Options and Parameters	148
TABLE 12-18	<code>reset</code> Command Options and Parameters	149
TABLE 12-19	<code>set</code> Command Options and Parameters	150
TABLE 12-20	<code>shutdown</code> Command Options and Parameters	152
TABLE 12-21	<code>sys</code> Command Options and Parameters	153
TABLE 12-22	<code>vol</code> Command Options and Parameters	155
TABLE 12-23	<code>volslice list</code> Command Options and Parameters	160
TABLE 12-24	<code>volslice create</code> Command Options and Parameters	162
TABLE 12-25	<code>volslice remove</code> Command Options and Parameters	164
TABLE 12-26	<code>lun map list</code> Command Options and Parameters	166
TABLE 12-27	<code>lun map add</code> Command Options and Parameters	168
TABLE 12-28	<code>lun map rm</code> Command Options and Parameters	169
TABLE 12-29	<code>lun perm list</code> Command Options and Parameters	172
TABLE 12-30	<code>lun perm</code> Command Options and Parameters	174
TABLE 12-31	<code>lun default</code> Command Options and Parameters	176
TABLE 12-32	<code>lun wwn rm</code> Command Options and Parameters	178
TABLE 12-33	<code>hwwn list</code> Command Options and Parameters	180
TABLE 12-34	<code>hwwn add</code> Command Options and Parameters	182
TABLE 12-35	<code>hwwn rm</code> Command Options and Parameters	183
TABLE 12-36	<code>hwwn rmgrp</code> Command Options and Parameters	184
TABLE 12-37	<code>hwwn add</code> Command Options and Parameters	186
TABLE 12-38	Topology Inquiry Results for SANbox Switch	187
TABLE 12-39	Topology Inquiry Results for Brocade Switch	187
TABLE 12-40	<code>set timezone</code> Command Options and Parameters	190
TABLE 12-41	Olson Time Zone Regions	190

TABLE 12-42	<code>ntp poll</code>	Command Options and Parameters	198
TABLE 12-43	<code>ntp interval</code>	Command Options and Parameters	199
TABLE 12-44	<code>ntp server</code>	Command Options and Parameters	200
TABLE 13-1		Message Types	202
TABLE 13-2		FRU Identifiers	203
TABLE 13-3		Error Messages	203
TABLE 13-4		Warning Messages	208
TABLE 13-5		Reset Log Types	219
TABLE 13-6		CLI Error Message Types	220
TABLE 13-7		Volume-Related (VN) Errors	221
TABLE 13-8		Port Errors	225
TABLE 13-9		Unit-Related Errors (Interconnect Card and Other FRUs)	226
TABLE 13-10		Embedded Operating System and Driver Errors	229

Preface

The *Sun StorEdge T3+ Array Administrator's Manual* describes how to configure, monitor, and troubleshoot Sun StorEdge™ T3+ arrays using the command-line interface (CLI). It provides information on commands that are specific to the array, but it is not intended as a reference for general operating system commands.

This guide is designed for use with the *Sun StorEdge T3+ Array Installation and Configuration Manual* and is written for an experienced system administrator of the Solaris™ operating environment and related disk storage systems.

Topics covered here are:

- “Before You Read This Book” on page xvi
- “Typographic Conventions” on page xvi
- “Related Documentation” on page xvii
- “Using UNIX Commands” on page xviii
- “Shell Prompts” on page xviii
- “Accessing Sun Documentation Online” on page xviii
- “Ordering Sun Documentation” on page xix
- “Sun Welcomes Your Comments” on page xix

Before You Read This Book

Read the *Sun StorEdge T3+ Array Installation, Operation, and Service Manual* to install the array.

Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	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.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Related Documentation

Application	Title	Part Number
Latest array updates	<i>Sun StorEdge T3+ Array Release Notes</i>	816-4771
Installation overview	<i>Sun StorEdge T3+ Array Start Here</i>	816-4768
Safety procedures	<i>Sun StorEdge T3+ Array Regulatory and Safety Compliance Manual</i>	816-0774
Site preparation	<i>Sun StorEdge T3+ Array Site Preparation Guide</i>	816-0778
Installation and configuration	<i>Sun StorEdge T3+ Array Installation and Configuration Manual</i>	816-4769
Administration	<i>Sun StorEdge T3+ Array Administrator's Manual</i>	816-4770
Cabinet installation	<i>Sun StorEdge T3 Array Cabinet Installation Guide</i>	806-7979
Disk drive specifications	<i>18 Gbyte, 1-inch, 10K rpm Disk Drive Specifications</i>	806-1493
	<i>36 Gbyte, 10K rpm 1-Inch Disk Drive Specifications</i>	806-6383
	<i>73 Gbyte, 10K rpm 1.6-Inch Disk Drive Specifications</i>	806-4800
Host Bus Adapters	<i>Sun Enterprise 6x00/5x00/4x00/3x00 Systems SBus+ and Graphics+ I/O Boards Installation Guide</i>	805-2704
	<i>Sun StorEdge PCI FC-100 Host Adapter Installation</i>	805-3682
	<i>Sun StorEdge SBus FC-100 Host Adapter Installation and Service Manual</i>	802-7572
	<i>Sun StorEdge PCI Single Fibre Channel Network Adapter Installation Guide</i>	806-7532
	<i>Sun StorEdge PCI Dual Fibre Channel Host Adapter Installation Guide</i>	806-4199
	<i>Sun StorEdge Compact PCI Dual Fibre Channel Network Adapter Installation and User's Guide</i>	816-0241

Using UNIX Commands

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

See one or more of the following for this information:

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

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	#

Accessing Sun Documentation Online

A broad selection of Sun™ system documentation is located at:

<http://www.sun.com/products-n-solutions/hardware/docs>

A complete set of Solaris documentation and many other titles are located at:

<http://docs.sun.com>

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

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

Sun Welcomes Your Comments

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

docfeedback@sun.com

Please include the part number (816-4770-10) of your document in the subject line of your email.

Overview of the Sun StorEdge T3+ Array

This chapter provides a high-level introduction to the Sun StorEdge T3+ array.

Topics in this chapter include:

- “Product Description” on page 2
- “Array Components” on page 4
- “Array Features” on page 11
- “Array Architecture” on page 12

Product Description

The Sun StorEdge T3+ array (see FIGURE 1-1) is a high-performance, modular, scalable storage device that contains an internal RAID controller and nine disk drives with Fibre Channel connectivity to the data host. Extensive reliability, availability, and serviceability (RAS) features include redundant components, notification of failed components, and the ability to replace components while the unit is online. The Sun StorEdge T3+ array provides the same features as the Sun StorEdge T3 array, and includes an updated controller card with direct fiber-optic connectivity and additional memory for data cache. The controller cards of both array models are described in more detail later in this chapter.

The array can be used either as a standalone storage unit or as a building block, interconnected with other arrays of the same type and configured in various ways to provide a storage solution optimized to the host application. The array can be placed on a table top or rackmounted in a server cabinet or expansion cabinet.

In this document, the Sun StorEdge T3+ array is referred to as the *array*.

The following characteristics of the array are provided:

- The array is a Fibre Channel disk storage solution
- Capability for up to nine 73-GByte drives in each array
- Supports RAID levels 0, 1, and 5
- Has up to 1 Gbyte of data cache memory
- Has hot-plugable FRUs

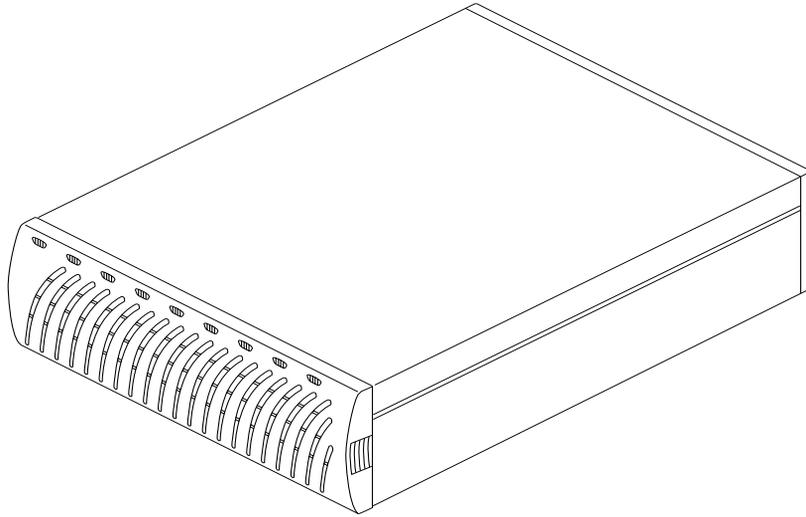


FIGURE 1-1 Sun StorEdge T3+ Array

Array Components

The array contains four basic components that can be easily replaced:

- Disk drive
- Controller card
- Interconnect card
- Power and cooling unit

All components plug into a centerplane; there is no internal cabling. For information on how to remove and replace these components, see the chapter on service in the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Note – The array centerplane and external chassis are physically connected, and are available as one field-replaceable unit (FRU). Only a qualified field-service representative should replace this FRU.

Disk Drive

Behind the panel at the front of the array are nine Fibre Channel disk drives, numbered drive 1 through drive 9 from left to right. See FIGURE 1-2. Each drive is in an enclosed canister that is easily installed and removed from the array. Drive *light-emitting-diodes* (LEDs), which are visible through the front panel, indicate drive activity and status. See *Sun StorEdge T3+ Array Installation and Configuration Manual*, for more information on LED locations and descriptions and how to service the disk drives.

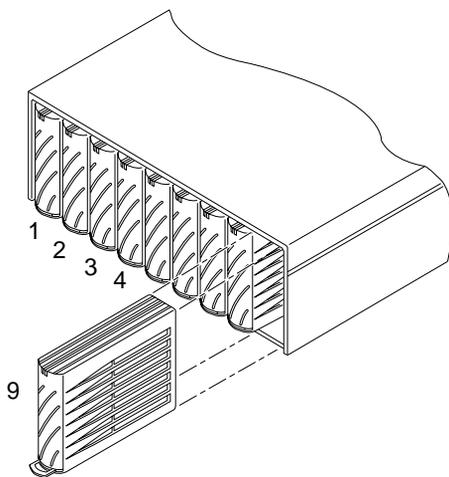


FIGURE 1-2 Disk Drives (Front View)

Array Controller Card

The Sun StorEdge T3+ array has a specific controller card that contains slightly different features on the exterior and a different design on the internal controller board. The general features of each array controller card are described in this section. See FIGURE 1-3.

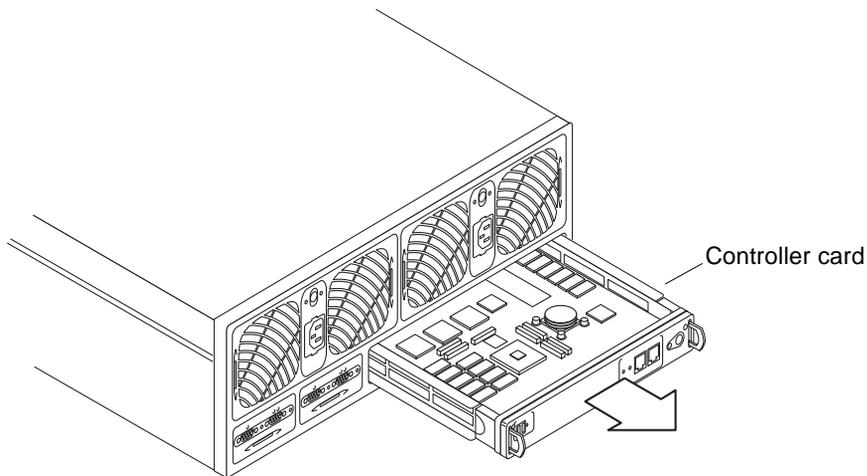


FIGURE 1-3 Controller Card (Rear View)

The location of the controller card for both models is shown in FIGURE 1-4. The *Sun StorEdge T3+ Array Installation and Configuration Manual*, provides information on how to service the controller card.

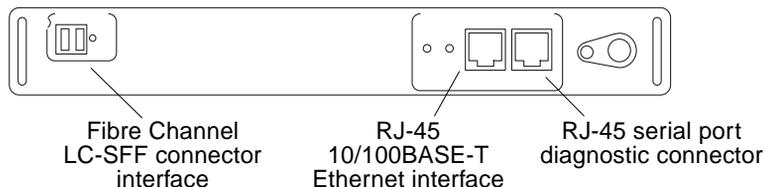


FIGURE 1-4 Sun StorEdge T3+ Array Controller Card Front Panel

The Sun StorEdge T3+ array controller card contains the RAID controller hardware and firmware, including:

- One Fibre Channel-Arbitrated Loop (FC-AL) port using an LC small-form factor (SFF) connector. The fiber-optic cable that provides data channel connectivity to the array has an LC-SFF connector that attaches directly to the port on the controller card. The other end of the fiber-optic cable has a standard connector (SC) that attaches a host bus adapter (HBA), hub, or switch.

- One 10/100BASE-T Ethernet host interface port (RJ-45). This port provides the interface between the controller card and the management host system. A shielded Ethernet cable (category 5) connects the controller to the site's network hub. This interface enables the administration and management of the array by way of the the command-line interface (CLI).
- One RJ-45 serial port. This serial port is reserved for diagnostic procedures that can only be performed by qualified service personnel.
- 1 Gbyte of high-speed SDRAM data cache.

FIGURE 1-5 shows the Sun StorEdge T3+ array controller card and connector ports.

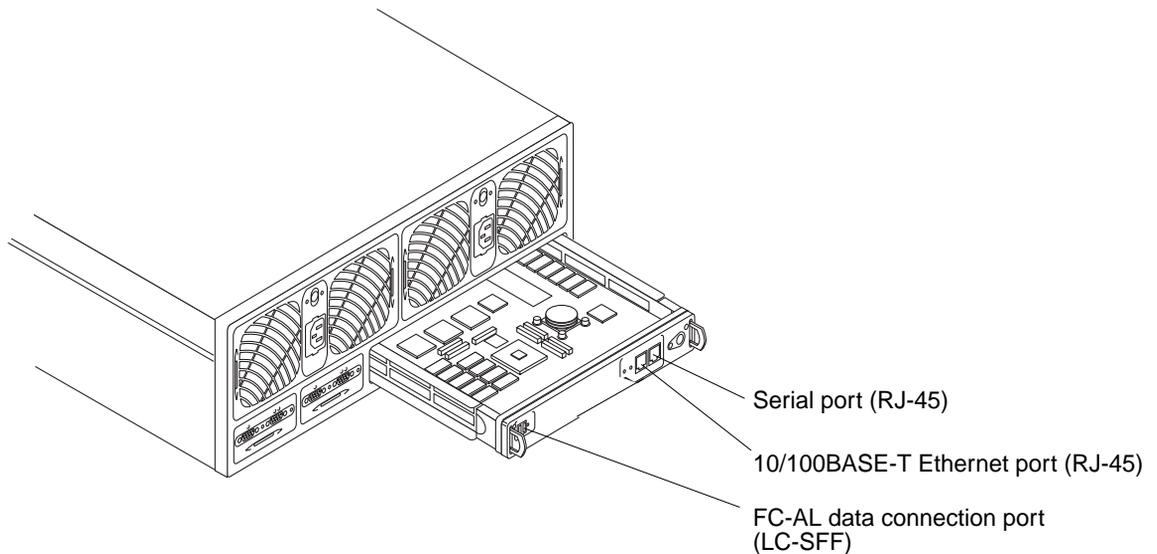


FIGURE 1-5 Sun StorEdge T3+ Array Controller Card and Ports

Unit Interconnect Cards (UIC)

The interconnect cards are alike on both array models. Each card has two interconnect ports: one input and one output for interconnecting multiple arrays.

The interconnect card provides switch and failover capabilities, as well as an environmental monitor for the array. Each array contains two interconnect cards for redundancy (thus providing a total of four interconnect ports).

FIGURE 1-6 shows the interconnect cards in a Sun StorEdge T3+ array.

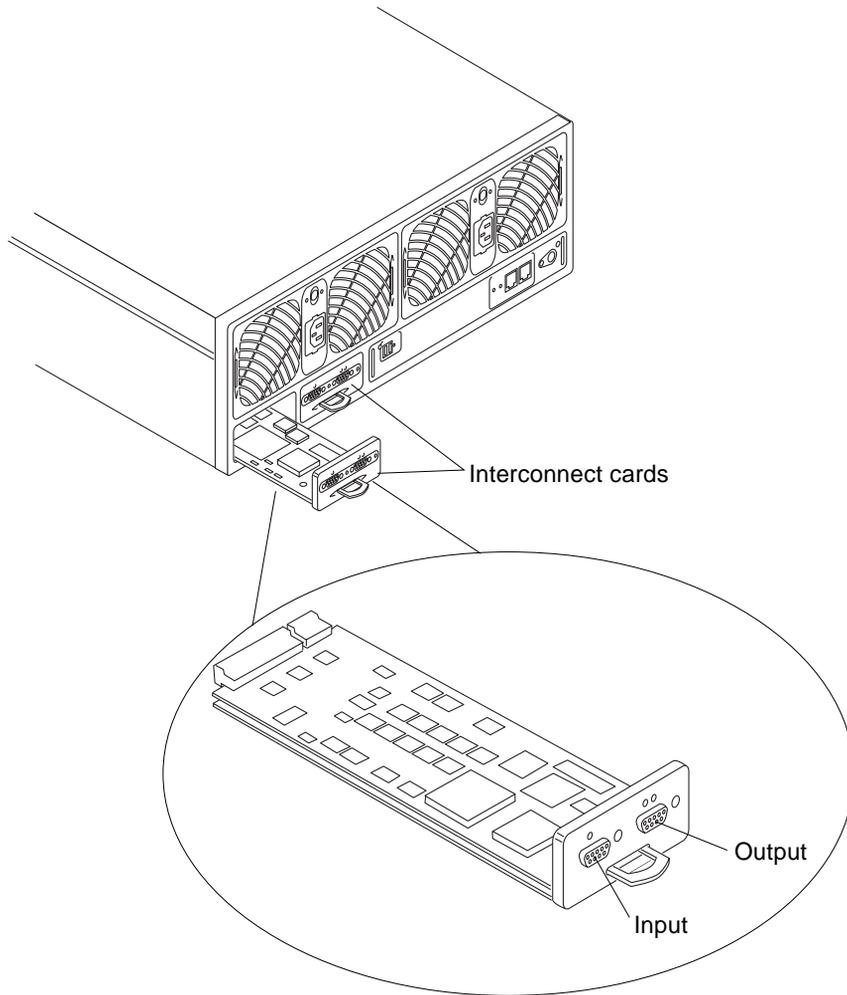


FIGURE 1-6 Interconnect Card and Ports

The interconnect card contains the interface circuitry and two connectors for interconnecting multiple arrays. It provides loop-switch capability and contains an environmental monitor for the array. Each array contains two interconnect cards for redundancy. For information on how to service the interconnect card, see *Sun StorEdge T3+ Array Installation and Configuration Manual*. See FIGURE 1-7.

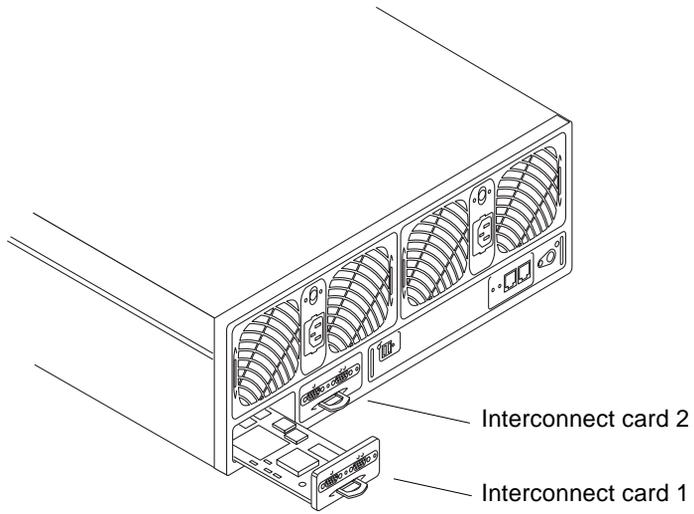


FIGURE 1-7 Interconnect Cards (Rear View)

Power and Cooling Unit

The power and cooling unit contains a power supply, two cooling fans, an integrated *uninterruptible power source* (UPS) battery, and LED status indicators for AC power. See FIGURE 1-8. Each array contains two power and cooling units for redundancy. See *Sun StorEdge T3+ Array Installation and Configuration Manual* for information on how to service the power and cooling unit and for information on UPS battery maintenance.

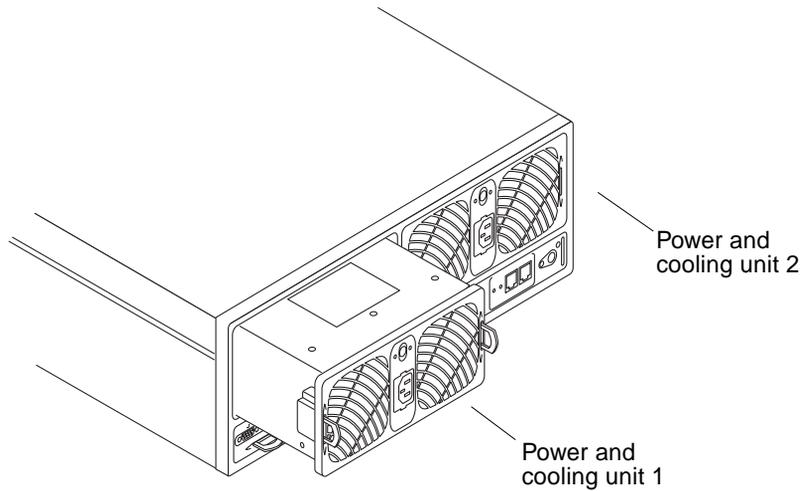


FIGURE 1-8 Power and Cooling Units (Rear View)

Array Features

The Sun StorEdge T3+ array provides high reliability while being simple to administer and maintain. The array features a design that supports high availability, high performance, configuration flexibility, and scalability of bandwidth, capacity, and input/output operations per second (IOPS). TABLE 1-1 provides an example of these key features, depicting an array with nine 73-Gbyte drives.

Note – Other drive capacities are available. Check with your Sun representative for more information on storage capacities that are currently available.

TABLE 1-1 Sun StorEdge T3+ Array Key Features—73-Gbyte Drive

Feature	Sun StorEdge T3+ Array
Capacity per unit	1.629 TB (with 9 x 181 GB drives) 1 GB cache
Performance per controller	95+ MB/sec for large RAID 5 reads from disk (64K >) 90+ MB/sec for large RAID 5 writes to disk 7,200 IOPS for 8 KB (100% cache hit reads)
Connectivity per controller	FC-AL host interface (LC-SFF to SC) 10/100BASE-T Ethernet network port
Scalability (one to two units, one to two controllers)	1,314 GB maximum capacity 1- to 2-GB cache 1 to 2 host Fibre Channel interfaces 90- to 190-MB/sec bandwidth 7,200 to 14,360 IOPS (100% cache hit reads)

TABLE 1-1 Sun StorEdge T3+ Array Key Features—73-Gbyte Drive (Continued)

Feature	Sun StorEdge T3+ Array
Reliability/redundancy	RAID 0/1 (1+0)/5
	Redundant back-end data loops
	Redundant interconnect cards
	Passive centerplane connector board
	Redundant power and cooling units and power cables
	Redundant back-end FC-AL interconnect cables
	Redundant controller configuration
Administration	Telnet session access (CLI interface)
	SNMP Version 1

Array Architecture

The Sun StorEdge T3+ array is uniquely designed to be modular, scalable, reliable, serviceable, high-performance storage building blocks with flexible configurations. The design allows for multiple arrays to be combined in various ways to provide complete storage solutions that are optimized for all applications—from transaction processing to decision support and high-performance computing, from workgroup environments to data centers.

The arrays have redundant components for high reliability. Components are hot swappable and field replaceable for serviceability. Enterprise configurations support host-based alternate pathing for enhanced availability.

Data and administrative paths are completely independent for reliability, security, serviceability, and ease of use. The network-based administrative path allows for centralized configuration and monitoring of large numbers of Sun StorEdge T3+ array configurations, providing storage to multiple application servers.

Each disk drive has a drive label, a small portion of which is reserved for the *system area*. Approximately 150 Mbytes is reserved for the system area, which contains the configuration data, boot firmware, and file system information. The system area information is mirrored across all nine drives for redundancy so that data can be recovered from the other functional drives.

Finally, the unique switched-loop architecture provides the flexibility to configure multiple units together for scalability and availability, while enabling loops to be reconfigured dynamically for diagnosis and recovery in case of loop-related failures.

The following is a summary of the characteristics of the array:

- Modular building blocks of disk storage
- Switched-loop architecture
- Independent data and administrative paths
- 100-MByte per second Fibre Channel data path
- Administrative path is networked-based and serial-based (RS-232)
- Two separate serial buses supporting control and sense operations.

Array Firmware 2.1 New Features Concepts

The Sun StorEdge T3 Array Controller Firmware version 2.1 provides new features, as described in this chapter.

Topics in this chapter include:

- “Firmware Version 2.1 New Features List” on page 16
- “Firmware Version 2.1 New Features Concepts” on page 16

Firmware Version 2.1 New Features List

The following is a complete list of all firmware version 2.1 new features:

- “Volume Slicing, Mapping, and Masking” on page 16
- “Auto Upgrade” on page 20
- “Split Loop Mode” on page 20
- “Fabric Point-to-Point Support” on page 20
- “Persistent Group Reservation” on page 20
- “Network Time Protocol” on page 21
- “Auto-Versioning” on page 21

Firmware Version 2.1 New Features Concepts

Volume Slicing, Mapping, and Masking

A selected number of disks within a Sun StorEdge T3+ array can be logically grouped together to form a volume. Each array unit can have no more than two volumes (volume 0 or volumes 0 and 1 by default—see FIGURE 2-1). Each volume is seen by a host as a *logical unit number* (LUN). Thus, a single volume in an array is said to be *mapped* to the host as a single LUN (1:1). When you upgrade to firmware version 2.1, nothing changes with respect to the volumes or LUNs, or with their 1:1 mapping (see FIGURE 2-2). Upon upgrading to firmware version 2.1, you can issue the `sys enable_volslice` command that enables volume slicing and automatically defines a default *slice* in software for each existing volume (with 1:1 mapping and having the same size as the volume) in an array. See FIGURE 2-3.

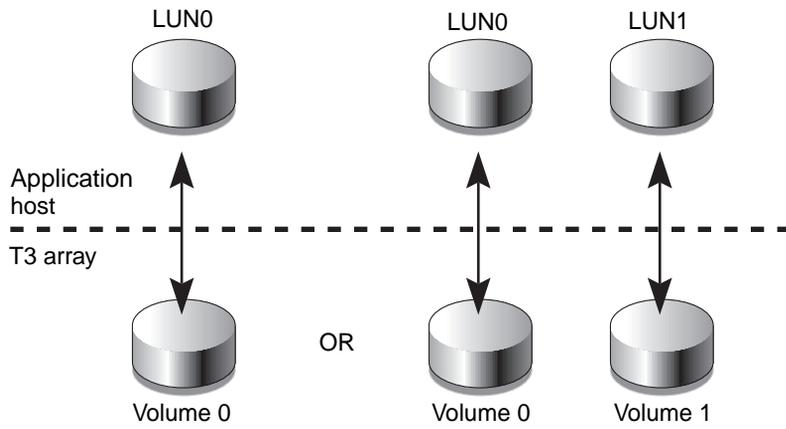


FIGURE 2-1 Two Possible Mapping Configurations for a Single T3+ Array Unit With Controller Firmware Version 2.0.1

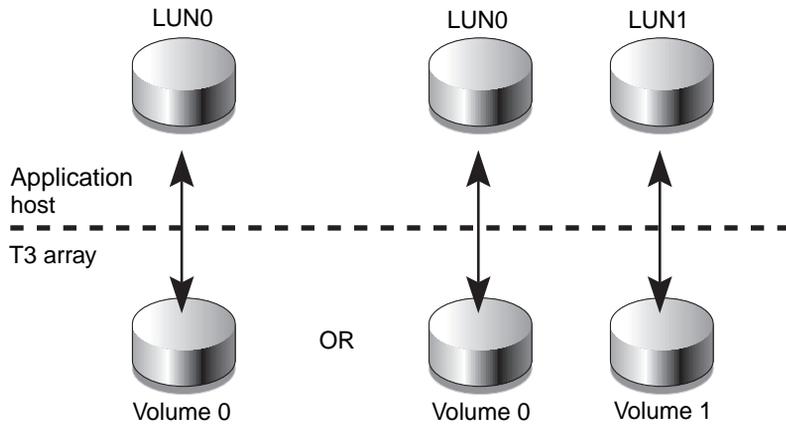


FIGURE 2-2 Two Possible Mapping Configurations for a Single T3+ Array Unit with Controller Firmware Version 2.1

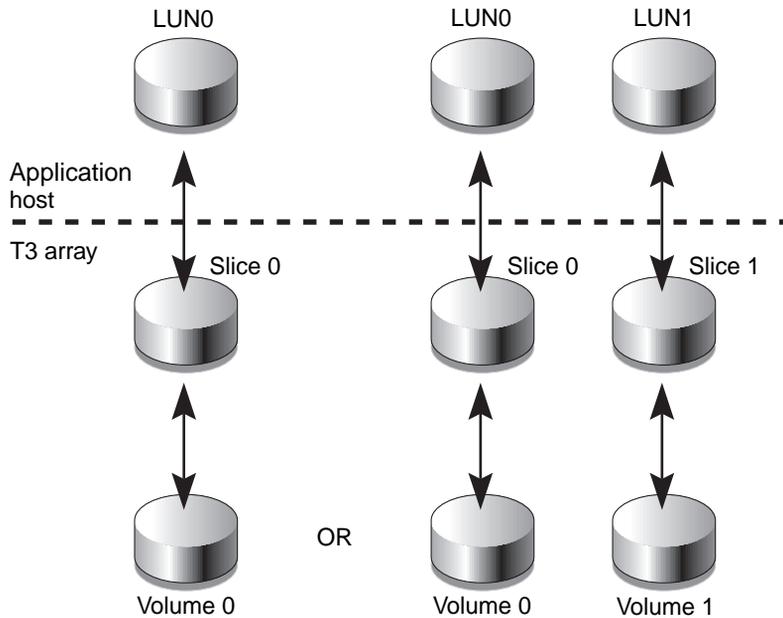


FIGURE 2-3 Two Possible Mapping Configurations for a Single T3+ Array Unit With Controller Firmware Version 2.1 and Volume Slicing Enabled

Firmware version 2.1 allows for multiple volume slicing (more than one slice per volume) which is similar to partitioning a hard drive of a desktop computer. This new feature allows greater flexibility to address storage in smaller increments, take advantage of larger capacity drives, and provides ease of scaling to suit individual applications. A 20-Gbyte volume can be configured in the following ways:

- Ten 2-Gbyte slices/LUNs
- Five 4-Gbyte slices/LUNs

A maximum of 16 slices can be made for an array workgroup or an enterprise configuration. Each slice is then mapped to a LUN with the corresponding maximum of 16 LUNs per workgroup or enterprise configuration. The maximum number of volumes per array unit remains at two. See FIGURE 2-4.

LUN masking is the term used for assigning LUN permissions (read-only, read/write, or none) to a host.

LUN masking is a key component of building large storage systems. It eases storage administration while allowing for a more secure environment. When a LUN is masked from a host it is not available to be configured and thus prevents the LUN

from being assigned to multiple hosts accidentally. This becomes critical for large network configurations. In such networks, if each host could access all of the LUNs, it would be difficult to administratively manage the host permissions.

Each host HBA port is assigned a unique world-wide number (WWN). LUN masking requires the verification each host I/O with an ID by checking the host port WWN against the permissions. This results in significant system security benefits.

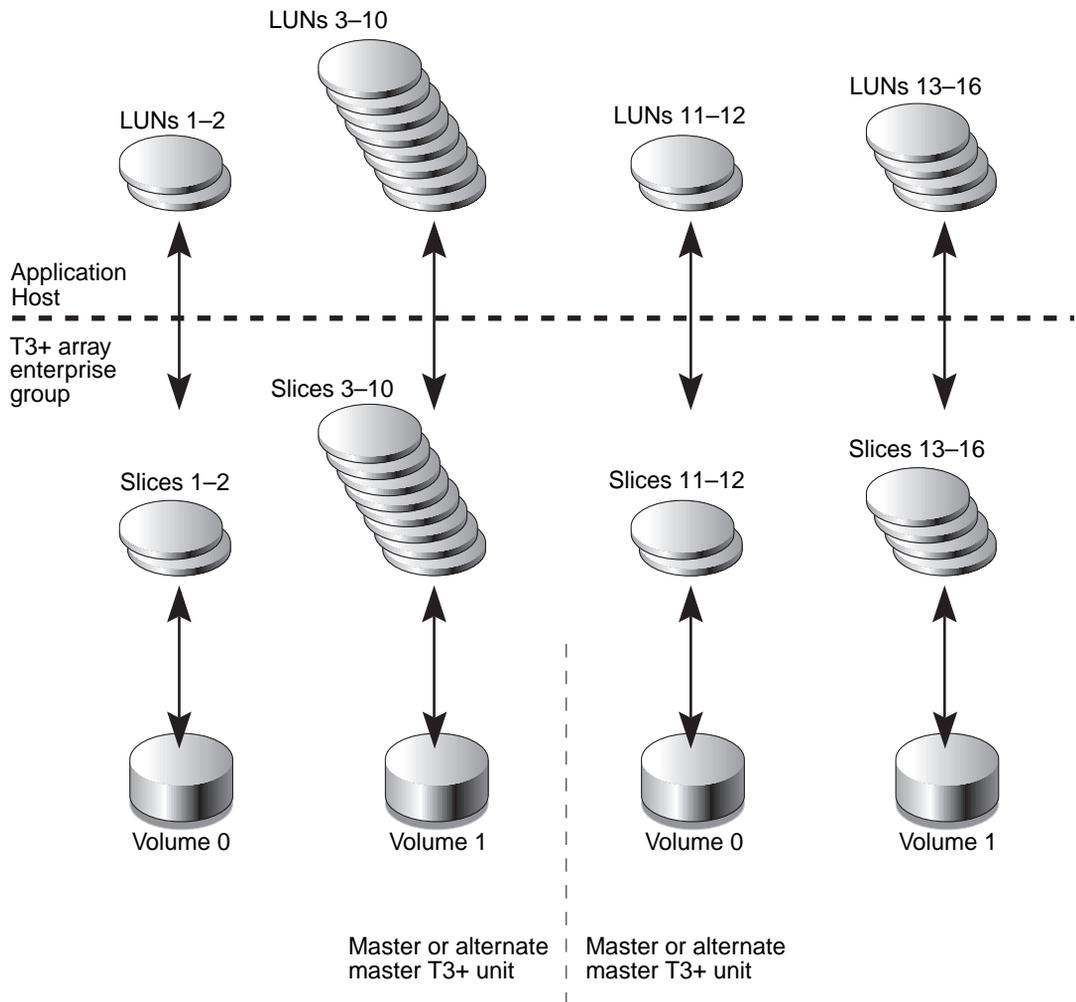


FIGURE 2-4 A Slicing Example for a T3+ Array Enterprise Group With Each T3+ Unit Having Two Volumes, Controller Firmware Version 2.1, and Volume Slicing Enabled

Auto Upgrade

This feature is part of the controller 2.1 firmware and is disabled by default. It can be activated with the `sys enable_volslice` command. The command must be issued to both arrays in an enterprise configuration. The command enables volume slicing and masking as described in “Volume Slicing, Mapping, and Masking” on page 16. See FIGURE 2-4. Auto Upgrade also performs the following functions:

- Sets up a default configuration of LUN Mapping
- Maintains WWNs of existing LUNs
- Preserves pre-2.1 data through the upgrade

Split Loop Mode

Split loop mode is a new feature built into the controller 2.0.1 firmware release that allows multiple back-end loops to the physical drives. This increases performance by utilizing the available bandwidth of the Fibre Channel loops. Formerly, the maximum 18 drives of an enterprise configuration could be accessed through each Interconnect card (18 through Fibre Channel loop A and 18 through loop B). The new split loop software on the controller now recognizes Fibre Channel loop A as being split into two loops. This allows 6 drives to communicate through Fibre Channel loop A1, 6 through loop A2, and 6 through loop B.

Fabric Point-to-Point Support

Prior versions of controller firmware only supported a Fibre Channel loop topology connectivity. With the 2.1 firmware release, a point-to-point (p2p) topology is now also supported. This means that a Sun StorEdge T3+ array can now connect directly to a switch, hub, or host. And, obviously, no arbitration is required for point-to-point communication.

Persistent Group Reservation

Persistent group reservation (PGR) is a new feature built into the controller 2.0.1 firmware release that implements the SCSI-3 Persistent Reserve commands on a Sun StoreEdge T3+ array. This feature can be used by any software that supports industry standard SCSI-3 Persistent Reserve commands.

Network Time Protocol

The Network Time Protocol (NTP) support feature allows a network system of Sun StorEdge T3+ arrays to automatically synchronize their internal system time with a designated NTP master Sun StorEdge T3+ array.

NTP is a highly advanced and complex distributed network clock synchronization protocol and is capable of achieving low-millisecond accuracy. NTP synchronizes and converts the T3+ array system clock with the Coordinated Universal Time (UTC, also known as Greenwich Mean Time (GMT)). A time zone module within UTC converts UTC time to local time and also checks and applies the local daylight saving rules and inserts seconds periodically to account for leap years.

With future releases of Solaris, Olson time zone names (see TABLE 12-41) will be adopted in addition to the traditional Solaris time zone names. This time zone standard is being adopted in many industries.

Auto-Versioning

Auto-versioning is a new feature built into the controller 2.0.1 firmware release that allows a Sun StorEdge T3 controller to be seamlessly upgraded with a Sun StorEdge T3+ controller without any impact resulting from the different booting processes of the two arrays.

Array Administration Overview

This book provides information on using the command-line interface (CLI) of the Sun StorEdge T3+ array for administrative tasks. Various chapters list CLI commands as they pertain to performing various tasks related to configuring, monitoring or troubleshooting an array.

Note – Throughout this book, the terms *logical volume*, *LUNs* (logical units), and *disks* are used interchangeably.

Topics in this chapter include:

- “To Display Commands” on page 24
- “To Display Command Syntax” on page 24
- “FRU Identifiers” on page 25

▼ To Display Commands

- To view the available commands on the array, type `help` at the prompt.

```
:/: help
arp      cat      cd      cmp      cp      date     echo     head
help     ls       mkdir   mv       ping    pwd      rm       rmdir
tail     touch
boot     disable disk    enable  fru     id       logger  lpc
more     passwd  port    proc     reset   set      shutdown
sync     sys     tzset   ver      vol     ep       refresh route
ofdg lun      hwn
```

Note – Because this book is not a reference manual for the operating system, not all of the commands displayed, such as common UNIX commands, are discussed in this book. See Chapter 12 for detail on commands specific to the array.

▼ To Display Command Syntax

- Type `command-name help` on the array. To display the syntax of a command,

For example:

To view the syntax of the `sys` command, type `sys help`.

```
:/: sys help
usage:  sys list
        sys stat
        sys blocksize <16k | 32k | 64k>
        sys cache <auto | writebehind | writethrough | off>
        sys mirror <auto | off>
        sys mp_support <none | rw | mpzio | std>
        sys naca <off | on>
        sys rd_ahead <on | off>
        sys recon_rate <high | med | low>
        sys enable_volslice
        sys fc_topology <auto | loop | fabric_p2p>
```

FRU Identifiers

Many commands use a field-replaceable unit (FRU) identifier to refer to a particular FRU in an array. This identifier contains a unit constant (*u*), sometimes referred to as the *enclosure*; the unit number (*n*); the FRU constant (*ctr* for controller card, *pcu* for power and cooling unit, *i* for interconnect card, *d* for disk drive); and the FRU number (*n*). See TABLE 3-1.

TABLE 3-1 FRU Identifiers

FRU	Identifier	Unit number
Controller card	<i>unctr</i>	<i>n</i> = unit number (1, 2, ...)
Power and cooling unit	<i>unpcun</i>	<i>n</i> = unit number (1, 2, ...) <i>n</i> = pcu number (1, 2)
Interconnect card	<i>unln</i>	<i>n</i> = unit number (1, 2, ...) <i>n</i> = interconnect card number (1, 2)
Disk drive	<i>undn</i>	<i>n</i> = unit number (1, 2, ...) <i>n</i> = disk drive number (1, 2, ... 9)

Resetting Global Parameters

This chapter contains information on modifying array settings within an *administrative domain*. An administrative domain can be either a single array controller unit (workgroup configuration) or a partner group (enterprise configuration) that shares common administration through a master controller. Change these settings to configure your array to suit your application and recovery needs.

Note – Cache segment size must be set before volumes are created. Thereafter, volumes must be removed before the cache segment block size can be changed.

Topics in this chapter include:

- “To Set the Cache Block Size” on page 28
- “To Enable Partner Group Multipathing” on page 30
- “To Set the Cache Mode” on page 30
- “to Enable Mirrored Cache” on page 31
- “To Set the LUN Reconstruction Rate” on page 32
- “To Perform Volume Verification” on page 33
- “To Set Cache Read-Ahead Threshold” on page 34
- “To Set the IP Address” on page 35

▼ To Set the Cache Block Size



Caution – Changing the cache segment block size requires that you delete the existing volume; deleting a volume will destroy your data. Back up all data before beginning this procedure.

The *data block size* is the amount of data written to each drive when striping data across drives. (The block size is also known as the *stripe unit size*.) The block size can be changed only when no volumes are defined. The block size can be configured as 16 Kbytes, 32 Kbytes, or 64 Kbytes. The default block size is 64 Kbytes.

A cache segment is the amount of data being read into cache. A cache segment is 1/8 of a data block. Therefore, cache segments can be 2 Kbytes, 4 Kbytes, or 8 Kbytes. Because the default block size is 64 Kbytes, the default cache segment size is 8 Kbytes.

Cache block size is universal throughout an administrative domain. Therefore, you cannot change it after you have created a volume. The only way to change it is to delete a volume, change the block size, and then create a new volume:

1. From the array, delete the existing volume as follows:

```
#!/: vol list

volume      capacity  raid  data    standby
v0          143.2 GB   5     u1d1-9  none
#!/: vol unmount volume-name

#!/: vol remove volume-name
```

2. Type `sys list` to determine the current cache segment size in blocks.

Cache segment size is shown as `blocksize`.

```
://: sys list
blocksize           : 64k
cache               : auto
mirror              : auto
mp_support           : none
naca                 : off
rd_ahead            : on
recon_rate          : med
sys memsize         : 128 MBytes
cache memsize       : 1024 MBytes
enable_volslice    : on
fc_topology         : auto
```

3. Specify and confirm the blocksize using the `sys` command.

a. Type `sys blocksize n` to change the blocksize, where $n = 16k, 32k,$ or $64k$.

b. Type `sys list` to display the revised blocksize.

For example:

```
://: sys blocksize 16k
://: sys list
blocksize           : 16k
cache               : auto
mirror              : auto
mp_support           : rw
naca                 : off
rd_ahead            : on
recon_rate          : med
sys memsize         : 128 MBytes
cache memsize       : 1024 MBytes
enable_volslice    : on
fc_topology         : auto
```

▼ To Enable Partner Group Multipathing

This task applies only when two controller units are configured as a partner group. For VERITAS Dynamic Multi-Pathing (DMP) by VERITAS Volume Manager, and the Sun StorEdge Traffic Manager Software (STMS) software to function properly, you must enable multipathing support on the array for enterprise configurations.

- **On the array, enable multipathing support by typing:**

- For DMP:

```
:/: sys mp_support rw
```

- For STMS:

```
:/: sys mp_support mpzio
```

Note – If you are using Sun StorEdge Traffic Manager software for your multipathing software, type `sys mp_support mpzio` to select the correct array setting to support this software.

▼ To Set the Cache Mode

Write caching improves performance for applications because the data is written to fast memory as opposed to slower disk storage devices. The cache is protected with redundant battery systems that ensure the data is written to disk in the event of an AC power loss.

Four possible cache mode options exist:

- Auto
- Write-behind
- Write-through
- Off

The default cache setting on the Sun StorEdge T3+ array is auto, which enables write-behind caching as long as there are no FRU failures. For more information about these settings, refer to Chapter 3 in the *Sun StorEdge T3+ Array Installation and Configuration Guide*.

- **Type** `sys cache cache-mode`.

For example, if you are setting the cache mode to write-behind, type:

```
:/: sys cache writebehind
```



Caution – Data loss can occur in the write-behind cache mode if the units are not configured as a fully redundant enterprise and an array system failure occurs. Use the auto cache mode for the highest degree of data protection.

Tip – You can view the cache status by typing `vol mode`.

▼ to Enable Mirrored Cache

Mirrored cache can safeguard cached data when a controller fails in an enterprise configuration. When enabled, cache data is copied to a cache mirror segment on the alternate controller before the write is acknowledged to the host.

Note – Cache mirroring impacts performance because the data is written to two separate controller caches.

- **Type:**

```
:/: sys mirror auto
```

Note – The `auto` setting enables cache mirroring only when array health and conditions are optimal.

▼ To Set the LUN Reconstruction Rate

You can configure the LUN reconstruction rate so application performance is not impacted. Reconstruction rate options are high, medium, and low.

Note – You cannot change the reconstruction rate while disk drive reconstruction is in progress.

1. Use the `sys list` command to check the current reconstruction rate (`recon_rate`).

```
:/: sys list
blocksize      : 16k
cache          : auto
mirror         : auto
mp_support     : rw
naca           : off
rd_ahead       : on
recon_rate     : med
sys memsize    : 128 MBytes
cache memsize  : 1024 MBytes
enable_volslice : on
fc_topology    : auto
```

2. Use the `sys` command to specify and confirm the reconstruction rate.

a. Type `sys recon_rate [high|med|low]` to change the reconstruction rate.

b. Type `sys list` to display the revised rate.

For example:

```
:// sys recon_rate low
:// sys list
blocksize           : 16k
cache               : auto
mirror              : auto
mp_support          : rw
naca                : off
rd_ahead            : on
recon_rate          : low
sys memsize         : 128 MBytes
cache memsize       : 1024 MBytes
enable_volslice    : on
fc_topology         : auto
```

▼ To Perform Volume Verification

The `vol verify` command executes manual parity checks on existing volumes. Parity checking applies only to RAID 1 and RAID 5 volumes. You check data parity using the `vol verify` command before performing tape-backup overwrite cycles, approximately once every 30 days.



Caution – Verify that system health is in optimal condition before you run the `vol verify` command. For example, ensure that no LUNs are under reconstruction; the status of all disks is zero (which can be checked using the `vol stat` command) and other similar conditions are resolved.



Caution – The `vol verify` operation can take up to several hours to run, depending on system activity and the verification rate selected. Executing this command can affect system performance, depending on these factors.

- Use the `vol verify` command to set parity check rate.

```
:/: vol verify volume-name [fix] rate n
```

where *volume-name* is the name of the volume to verify.

Note – The volume name is a name internal to the array and is not seen by the host.

- The `[fix]` option corrects parity errors on RAID 5 volumes and corrects mirrored data errors on RAID 1 volumes. If `fix` is *not* specified, then `vol verify` reports errors, but does not correct them. If the `fix` option is specified and an error is detected, the `vol verify` command regenerates parity from the existing data on the volume.
- The verification rate is *n*, where *n* equals any number from 1 to 8. The default rate is 1, which has the minimum performance impact on the data host.

Note – Within an enterprise configuration, the `vol verify` command can only be run on one volume at a time.

▼ To Set Cache Read-Ahead Threshold

Setting read caching can reduce disk I/O by reading data for future retrieval. This refers to the number of sequential reads in a chunk of data. The two cache read-ahead threshold options are `on` and `off`. The default setting for the array is `on`.

1. Use the `sys list` command to check the current cache read-ahead threshold (`rd_ahead`).

```
:/: sys list
blocksize      : 16k
cache          : auto
mirror        : auto
mp_support    : rw
naca          : off
rd_ahead      : on
recon_rate    : low
sys memsize   : 128 MBytes
cache memsize : 1024 MBytes
enable_volslice : on
fc_topology   : auto
```

2. Use the `sys rdAhead` command to set the default threshold to `off` and confirm the setting using the `sys list` command.

For example:

```
://: sys rdAhead off
://: sys list
blocksize           : 16k
cache               : auto
mirror             : auto
mp_support          : rw
naca               : off
rdAhead            : off
recon_rate         : low
sys memsize        : 128 MBytes
cache memsize      : 1024 MBytes
enable_volslice    : on
fc_topology        : auto
```

▼ To Set the IP Address

You can set the IP address on an array with either of the following methods:

- Edit the host files on a reverse address resolution protocol (RARP) server with IP address information.

This option enables the IP address to be downloaded automatically when the array is first installed and powered on, and then receives a RARP request from the host. If you move the array to another network, it receives its IP information from the RARP server on that network; thus, the host files of the new RARP server must be updated to include the new array. To set the array IP address by editing the host files, refer to the instructions in the Installation section of the *Sun StorEdge T3+ Array Installation and Sun StorEdge Configuration Guide*.

- Use the array `set ip` command.

This option can be used only if the array is online and already connected to an Ethernet network, which would be possible only if the array IP address had been assigned using the RARP server previously. Use this option if you are moving the array to a network without a RARP server.

- **Type `set ip IP-address` to assign the array IP network address:**

```
://: set ip IP-address
```



Caution – If you are moving the array to a network with a RARP server to assign a new IP address with this command, be aware that if the array reboots, it can pick up another IP address from the RARP server that will override the IP address you assigned with the `set ip` command. To avoid this, always assign the IP address by editing the host files on a RARP server.

Reconfiguring a Volume

This chapter contains information on configuring volume size, RAID level, and a hot spare.

Topics in this chapter include:

- “Reconfiguration Limitations” on page 38
- “To Delete a Logical Volume” on page 38
- “To Create a Logical Volume” on page 39
- “To Label a Volume” on page 41

Note – Cache block size must be configured before volumes are configured. See Chapter 5, “Resetting Global Parameters” on page 27 for more information.

Reconfiguration Limitations

Volumes can only be created or deleted. Once a volume has been configured, you cannot reconfigure it to change its size, RAID level, or hot-spare configuration. You can only delete the volume and create a new one with the desired configuration.



Caution – Before you begin deleting and creating new volume(s), back up all data on the array. Data that resides in the volume to be deleted will be destroyed.

▼ To Delete a Logical Volume



Caution – Deleting a volume will destroy your data. Back up all data before beginning this procedure.

1. Use the `vol list` command to display the configuration of the current volumes.

```
:/: vol list

volume      capacity  raid  data  standby
v0          143.2 GB   5     u1d1-9  none
```

2. Unmount the volume.

```
:/: vol unmount volume-name
```

3. Delete the volume.

```
:/: vol remove volume-name
```

▼ To Create a Logical Volume

When configuring a logical volume, you must define the volume, RAID level, and hot spare at the same time.

Note – Configuring a hot-spare drive in the volume is optional. If you are configuring a hot spare in the volume, you can use only disk drive 9 (d9) in the array as a hot spare. You can have only one hot spare per array, which can be shared between two volumes.

TABLE 5-1 shows examples of valid volume configurations.

TABLE 5-1 Volume Configuration Examples

Volume 1	Volume 2	Hot spare
9 disk RAID 5	None	
8 disk RAID 5	None	X
9 disk RAID 1	None	
8 disk RAID 1	None	X
2 disk RAID 1	7 disk RAID 5	
2 disk RAID 1	6 disk RAID 5	X
2 disk RAID 1	7 disk RAID 1	
2 disk RAID 1	6 disk RAID 1	X
4 disk RAID 0	5 disk RAID 1	
2 disk RAID 0	6 disk RAID 5	X
4 disk RAID 5	5 disk RAID 5	

1. On the array, use the `vol add` command to create the volume as follows:
 - a. Define the volume name (`vol add volume-name`).
 - b. Define the drives (`data undn-n`) on which the volume will reside, where:
 - `un` is the array unit number
 - `dn-n` are the disk drives, $n = 1$ to 9
 - c. Define the RAID level (`raid n`), where $n = 0, 1, \text{ or } 5$.

d. Optional: define the hot-spare drive (standby und9) where:

- un is the array unit number
- d9 is the number of the hot-spare disk drive

```
:/: vol add volume-name data undn-n raid n standby und9
```

For example:

```
:/: vol add v1 data u2d1-8 raid 5 standby u2d9
```

- v1 is the volume name
- u2d1-8 indicates the location of the volume: unit 2, disk drives 1 through 8
- raid 5 is RAID level 5
- standby u2d9 is the location of the hot spare: unit 2, drive 9

2. Check the status of the volumes.

The `vol stat` command shows the drive status. The status of all drives must be 0. For example:

```
:/: vol stat

v1          u2d1  u2d2  u2d3  u2d4  u2d5  u2d6  u2d7  u2d8  u2d9
unmounted  0      0      0      0      0      0      0      0      0
```

3. Use the `vol init` command to initialize the volume.

Depending on system activity at the time of initialization, it can take up to an hour to initialize a volume. Only one volume can be initialized at a time.

```
:/: vol init volume-name data
```

4. Use the `vol mount` command to mount the volume.

```
:/: vol mount volume-name
```

5. Use the `vol list` command to confirm that you created the volume correctly.

For example:

```
:/: vol list

volume          capacity  raid  data  standby
v1              125.2 GB   5    u2d1-8  u2d9
```

Note – If you are running on the Solaris 7, 11/99 operating environment or a later release of the Solaris operating environment, skip Step 6. Later versions of the Solaris operating environment automatically recognize added storage devices without additional command operations.

6. On the data host, use the `luxadm(1M)` command to recognize the new volume.

In the Solaris environment, the `luxadm(1M)` command probes for new devices. Refer to the `luxadm(1M)` man page for more information on this command.

On the host system, type `luxadm insert`:

```
# luxadm insert
```

Note – If the `luxadm` utility is not available, you will have to do a reconfiguration reboot (`reboot -- -r`) to ensure that the host recognizes the new volumes. It is preferable, however, to use the `luxadm` command for this procedure instead of the `reboot -- -r` command.

▼ To Label a Volume

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

To change the label of a volume:

1. On the array, type `vol stat` to verify that the volume is mounted.

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

Specify the disk number when prompted. In this example, the array is shown as disk number 2. It can be identified by the `SUN-T300-0200` label.

`format`

```
Searching for disks...done
```

AVAILABLE DISK SELECTIONS:

0. c0t2d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk8
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b230,0
1. c0t2d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk9
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b230,1
2. c0t4d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk10
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b37f,0
3. c0t4d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk11
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b37f,1
4. c0t6d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk12
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b28e,0
5. c0t6d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk13
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b28e,1
6. c0t8d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk14
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b719,0
7. c0t8d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk15
/pci@70,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b719,1
8. clt1d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk8
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000af55,0
9. clt1d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk9
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000af55,1
10. clt3d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk10
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000afdf,0
11. clt3d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk11
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000afdf,1
12. clt5d0 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk12
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b723,0
13. clt5d1 <SUN-T300-0200 cyl 34901 alt 2 hd 128 sec 256> disk13
/pci@71,2000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w50020f230000b723,1
14. c3t2d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
/sbus@7d,0/SUNW,fas@1,8800000/sd@2,0
15. c3t3d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
/sbus@7d,0/SUNW,fas@1,8800000/sd@3,0

```
Specify disk (enter its number): 2
```

```
selecting c0t4d0
```

```
[disk formatted]
```

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

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

```
FORMAT MENU:
  disk           - select a disk
  type           - select (define) a disk type
  partition      - select (define) a partition table
  current        - describe the current disk
  format         - format and analyze the disk
  repair         - repair a defective sector
  label          - write label to the disk
  analyze        - surface analysis
  defect         - defect list management
  backup         - search for backup labels
  verify         - read and display labels
  save           - save new disk/partition definitions
  inquiry        - show vendor, product and revision
  volname        - set 8-character volume name
  !<cmd>         - execute <cmd>, then return0
  quit
format> type
```

4. Type 0 to choose the Auto configure option.

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

```
AVAILABLE DRIVE TYPES:
  0. Auto configure
  1. Quantum ProDrive 80S
  2. Quantum ProDrive 105S
  3. CDC Wren IV 94171-344
  4. SUN0104
  5. SUN0207
  6. SUN0327
  7. SUN0340
  8. SUN0424
  9. SUN0535
 10. SUN0669
 11. SUN1.0G
 12. SUN1.05
 13. SUN1.3G
 14. SUN2.1G
 15. SUN2.9G
 16. SUN-T300-0100
 17. other
Specify disk type (enter its number)[16]: 0
clt1d0: configured with capacity of 33.34GB
<SUN-T300-0200 cyl 34145 alt 2 hd 16 sec 128>
selecting clt1d0
[disk formatted]
```

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

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

Monitoring the Array

This chapter describes how to use the CLI to monitor the Sun StorEdge T3+ array, and how to set up the array files for message logging and remote notification.

You can check array status by using a variety of CLI commands. This chapter discusses how to monitor the following:

- “To Determine Failover” on page 46
- “To Check the Drive Status” on page 47
- “To Check the Hot-Spare” on page 48
- “To Check Data Parity” on page 48
- “To Check the Battery” on page 49
- “To Display the FRU Information” on page 50
- “to Check the FRU Status” on page 51

▼ To Determine Failover

1. To determine which unit is the master or alternate master unit, type `sys stat`.
The following example shows a partner group in a normal state:

```
#!/: sys stat
Unit   State      Role    Partner
-----
 1     ONLINE     Master   2
 2     ONLINE     AlterM  1
```

In a failover state, unit 2 assumes the role of master unit and unit 1 is disabled, as shown in the following example:

```
#!/: sys stat
Unit   State      Role    Partner
-----
 1     DISABLED   Slave
 2     ONLINE     Master
```

2. To display how paths are mapped from the host ports to the volume, type `port listmap`.

```
#!/: port listmap

port   targetid  addr_type  lun  volume  owner  access
-----
ulp1   1          hard       0    v0      u1     primary
ulp1   1          hard       1    v1      u2     failover
u2p1   2          hard       0    v0      u1     failover
u2p1   2          hard       1    v1      u2     primary
```

▼ To Check the Drive Status

- Use the `vol stat` command to check drive status codes.

All drives should show a status of 0 under normal conditions.

```
:/: vol stat
```

```
v0          u1d1  u1d2  u1d3  u1d4  u1d5  u1d6  u1d7  u1d8  u1d9
mounted     0      0      0      0      0      0      0      0      0
v1          u2d1  u2d2  u2d3  u2d4  u2d5  u2d6  u2d7  u2d8  u2d9
mounted     0      0      0      0      0      0      0      0      0
```

The numeric drive status codes are listed in the following table.

TABLE 6-1 Drive Status Codes

Value	Description
0	The drive is mounted.
2	The drive is present.
3	The drive is spun up.
4	The drive is disabled.
5	The drive has been replaced.
7	An invalid system area exists on the drive.
9	The drive is not present.
D	The drive is disabled and is being reconstructed.
S	The drive is substituted.

▼ To Check the Hot-Spare

1. Use the `vol list` command to check the location of the hot spare (standby) drive.

```
:/: vol list

volume      capacity  raid  data      standby
v0          125.2 GB  5     u1d1-8   u1d9
v1          125.2 GB  5     u2d1-8   u2d9
```

2. Use the `vol stat` command to check the status of the hot spare drive.

```
:/: vol stat

v0          u1d1  u1d2  u1d3  u1d4  u1d5  u1d6  u1d7  u1d8  u1d9
mounted    0      0      0      0      0      0      0      0      0
```

All drives should show a status of 0. See TABLE 6-1 for definitions of drive status codes.

▼ To Check Data Parity

Note – Check data parity before performing tape-backup overwrite cycles, approximately once every 30 days.

- Use the `vol verify` command to perform a parity check of the drives.

```
:/: vol verify [fix] volume-name
```

where *volume-name* is the name of the volume to verify.

Note – The volume name is a name internal to the array and is not seen by the host.

- The [`fix`] option corrects parity errors on RAID 5 volumes and corrects mirrored data errors on RAID 1 volumes. If `fix` is *not* specified, then `vol verify` will report errors, but not correct them. If the `fix` option is specified and an error is detected, the `vol verify` command will regenerate parity from the existing data on the volume.

Note – The `vol` command is not re-entrant. Therefore, you cannot run other `vol` commands on the array until the `vol verify` operation has completed.

See “To Perform Volume Verification” on page 33 for more information on checking parity.

▼ To Check the Battery

1. Use the `id read` command to display information about battery life. (Unit number $n = 1$ or 2 ; power cooling unit number $n = 1$ or 2 .)

```
:/: id read unpcun
      Revision           : 0000
      Manufacture Week   : 00221999
      Battery Install Week: 00221999
      Battery Life Used  :  0 days, 0 hours
      Battery Life Span  : 730 days, 12 hours
      Serial Number      : 01204
      Vendor ID          : TECTROL
      Model ID           : 300-1454-01
```

2. Use the `refresh -s` command to check the status of a battery refresh cycle.

The following example shows a normal battery status (no refresh cycle):

```
:/: refresh -s

No battery refreshing Task is currently running.

      PCU1                      PCU2
      -----
      U1                        Normal                Normal
      U2                        Normal                Normal

      Current Time              Fri Jun 09 16:54:53 2000
      Last Refresh               Thu Jun 01 12:54:19 2000
      Next Refresh               Fri Jun 29 12:54:19 2000
```

▼ To Display the FRU Information

- Use the `fru list` command to display FRU vendor information, serial number, and firmware levels.

In the event of a FRU failure, `fru list` output is helpful in verifying correct FRU replacement because it contains the serial numbers. The following example shows `fru list` output for a Sun StorEdge T3+ array.

```
:/: fru list
```

ID	TYPE	VENDOR	MODEL	REVISION	SERIAL
ulctr	controller card	0034	501-5710-02(0200	123456
u2ctr	controller card	0034	501-5710-02(0200	123455
uld1	disk drive	SEAGATE	ST118202FSUN	9329	LK478728
uld2	disk drive	SEAGATE	ST118202FSUN	9329	LK493799
uld3	disk drive	SEAGATE	ST118202FSUN	9329	LK493800
uld4	disk drive	SEAGATE	ST118202FSUN	9329	LK494457
uld5	disk drive	SEAGATE	ST118202FSUN	9329	NK040486
uld6	disk drive	SEAGATE	ST118202FSUN	9329	LK339351
uld7	disk drive	SEAGATE	ST118202FSUN	9329	LK150715
uld8	disk drive	SEAGATE	ST118202FSUN	9329	NK040157
uld9	disk drive	SEAGATE	ST118202FSUN	9329	LKG79907
u2d1	disk drive	SEAGATE	ST118202FSUN	9329	LKG90019
u2d2	disk drive	SEAGATE	ST118202FSUN	9329	LKH18597
u2d3	disk drive	SEAGATE	ST118202FSUN	9329	LKH15606
u2d4	disk drive	SEAGATE	ST118202FSUN	9329	LKH16563
u2d5	disk drive	SEAGATE	ST118202FSUN	9329	LKG88883
u2d6	disk drive	SEAGATE	ST118202FSUN	9329	LKH61907
u2d7	disk drive	SEAGATE	ST118202FSUN	9329	LKG90719
u2d8	disk drive	SEAGATE	ST118202FSUN	9329	LKG95442
u2d9	disk drive	SEAGATE	ST118202FSUN	9329	LKG61774
u1l1	loop card	SCI-SJ	375-0085-01-	5.02 Flash	007924
u1l2	loop card	SCI-SJ	375-0085-01-	5.02 Flash	007382
u2l1	loop card	SCI-SJ	375-0085-01-	5.02 Flash	003343
u2l2	loop card	SCI-SJ	375-0085-01-	5.02 Flash	003344
ulpcu1	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001455
ulpcu2	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001408
u2pcu1	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001686
u2pcu2	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001445
ulmpn	mid plane	SCI-SJ	375-0084-01-	0000	000650
u2mpn	mid plane	SCI-SJ	375-0084-01-	0000	000649

▼ to Check the FRU Status

- Use the `fru stat` command to provide a status of each FRU.

Possible FRU states and definitions are listed in TABLE 6-2

```

:/: fru stat

```

CTLR	STATUS	STATE	ROLE	PARTNER	TEMP		
ulctr	ready	enabled	master	u2ctr	32.0		
u2ctr	ready	enabled	alt master	ulctr	29.0		

DISK	STATUS	STATE	ROLE	PORT1	PORT2	TEMP	VOLUME
uld1	ready	enabled	data disk	ready	ready	34	v0
uld2	ready	enabled	data disk	ready	ready	34	v0
uld3	ready	enabled	data disk	ready	ready	31	v0
uld4	ready	enabled	data disk	ready	ready	34	v0
uld5	ready	enabled	data disk	ready	ready	32	v0
uld6	ready	enabled	data disk	ready	ready	38	v0
uld7	ready	enabled	data disk	ready	ready	33	v0
uld8	ready	enabled	data disk	ready	ready	34	v0
uld9	ready	enabled	data disk	ready	ready	36	v0
u2d1	ready	enabled	data disk	ready	ready	30	v1
u2d2	ready	enabled	data disk	ready	ready	31	v1
u2d3	ready	enabled	data disk	ready	ready	34	v1
u2d4	ready	enabled	data disk	ready	ready	34	v1
u2d5	ready	enabled	data disk	ready	ready	32	v1
u2d6	ready	enabled	data disk	ready	ready	37	v1
u2d7	ready	enabled	data disk	ready	ready	38	v1
u2d8	ready	enabled	data disk	ready	ready	42	v1
u2d9	ready	enabled	data disk	ready	ready	34	v1

LOOP	STATUS	STATE	MODE	CABLE1	CABLE2	TEMP		
u2l1	ready	enabled	master	installed	-	30.0		
u2l2	ready	enabled	slave	installed	-	30.5		
u1l1	ready	enabled	master	-	installed	27.0		
u1l2	ready	enabled	slave	-	installed	29.5		

POWER	STATUS	STATE	SOURCE	OUTPUT	BATTERY	TEMP	FAN1	FAN2
ulpcu1	ready	enabled	line	normal	normal	normal	normal	normal
ulpcu2	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu2	ready	enabled	line	normal	normal	normal	normal	normal

TABLE 6-2 Possible FRU States and Definitions

FRU	Status	State	Definition
Controller Card			Status and State Conditions
	absent	disabled	Controller not in enclosure
	ready	disabled	Disabled or disabling
	offline	disabled	Reset or resetting
	booting	enabled	Booting
	ready	enabled	Online
Disk Drive			Status and State Conditions
	ready	enabled	Online
	fault	enabled	Problem with the drive; check LED
	missing	enabled	Drive is missing
	Role		Disk Drive Assignment
	data disk		Part of volume
	unassigned		Not in volume
	standby		Standby disk for volume. (Must be drive number 9)
	Port		Drive Port Condition¹
	ready		Corresponding drive port online
	not ready		Drive port not ready. Expect ready soon.
	bypass		This drive port is being bypassed; use another. (It is possible that the interconnect card with corresponding loop to drive is down.)
	unknown		Drive port state is not known. (Occurs only in partner groups when drive's corresponding interconnect card is down and interconnect card on partner unit is up.)

TABLE 6-2 Possible FRU States and Definitions *(Continued)*

FRU	Status	State	Definition
Interconnect (Loop) Card	Status and State Conditions		
	ready	enabled	Online
	missing		Interconnect card not installed
	offline		Interconnect card offline
	fault		Interconnect card has fault condition
Power and Cooling Unit	Status and State Conditions		
	ready	enabled	Online
	ready	substituted	Battery refresh in progress
	missing		Power and cooling unit not installed
	fault	enabled	Problem in power and cooling unit. For example: fan fault, output fault, over temperature.
	fault	disabled	Power and cooling unit offline
	Power Source		
	line		Powered from power cord
	battery		Powered from battery
	--		Power and cooling unit not installed
	Power and Cooling Unit Condition		
	normal		Online
	fault		Fault condition
	--		Power and cooling unit not installed
	Battery Condition		
	normal		Online
	fault		Battery refresh in progress
	--		Power and cooling unit not installed
	Temperature Condition		
	normal		Normal operating temperature
	fault		Power and cooling unit over temperature
	--		Power and cooling unit not installed

TABLE 6-2 Possible FRU States and Definitions *(Continued)*

FRU	Status	State	Definition
Power and Cooling Unit	Fan		Fan Condition
	normal		Online
	fault		Fault condition
	--		Power and cooling unit not installed

1. Each drive has two ports, and each port corresponds to an interconnect card. If one interconnect card fails, then all drives on that loop will bypass that particular port, and the interconnect card that is still online will start controlling all drives using the other port.

Array Configuration

Topics in this chapter include:

- “Configuring the Array for Remote System Logging” on page 56
- “Configuring the Array for SNMP Notification” on page 63

Configuring the Array for Remote System Logging

Configuring the array consists of editing the `/etc/syslog.conf` and the `/etc/hosts` file that exist on the array. Editing these files enables system messages to be logged and directed to a management host. Because you cannot edit files on the array, you must use FTP to transfer them to a host, make the edits on the host, and then transfer them back to the array.

This procedure consists of the following tasks:

- “To Transfer Array Files to the Management Host” on page 56
- “To Edit the Array `/etc/syslog.conf` File” on page 58
- “To Edit the Array `/etc/hosts` File” on page 59
- “To Transfer Files Back to the Array” on page 60
- “To Edit the Management Host `/etc/syslog.conf` File” on page 61

Note – For information on interpreting system messages, see Chapter 13.

▼ To Transfer Array Files to the Management Host

1. Start an `ftp` session from the management host to the array.

For example:

```
mngt_host: /:<15>ftp 129.146.81.201
Connected to 129.146.81.201.
220 chon-ji FTP server (SunOS 5.7) ready.
Name (129.146.81.201:root):
```

2. Log in to the array by typing `root` and your password at the prompts.

```
Name (129.146.81.201:root): root

331 Password required for root.
Password: password
230 User root logged in.
ftp>
```

3. Move to your working directory on the management host.

For example:

```
ftp> lcd /tmp
Local directory now /tmp
ftp>
```

4. Move to the /etc directory on the array.

```
ftp> cd /etc
250 CWD command successful.
ftp>
```

5. Type `binary` to set the transfer mode.

6. Copy the `syslog.conf` file from the /etc directory on the array to your working directory.

```
ftp> get syslog.conf
200 PORT command successful.
150 Binary data connection for syslog.conf (129.150.47.101,1031) (162 bytes).
226 Binary Transfer complete.
162 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

7. Copy the `hosts` file from the /etc directory on the array to your working directory.

```
ftp> get hosts
200 PORT command successful.
150 Binary data connection for hosts (129.150.47.101,1034) (47 bytes).
226 Binary Transfer complete.
47 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

8. Exit the ftp session.

```
ftp>
quit
221 Goodbye.
mngt_host:/:
```

▼ To Edit the Array `/etc/syslog.conf` File

You must edit the `/etc/syslog.conf` file to include the category of messages that you want to have logged with the IP address and host name of the management host that will be logging these messages.

- 1. Determine the IP address and host name of the management host that will be logging messages.**

Consult with your system administrator if you do not have this information.

- 2. Determine which category of messages you would like to receive from the array.**

Messages generated by the array are grouped into four categories in the order of severity, as described in the following table.

TABLE 7-1 Message Categories

Category	Description
Error	Indicates a critical system event requiring immediate user intervention or attention. For example, inability to flush the write-behind cache.
Warning	Indicates a serious system event requiring eventual user intervention. For example, a disk drive being disabled.
Notice	Indicates a system event that may lead to a more serious condition in the future. For example, the occurrence of hard errors corrected via parity replacement.
Information	Indicates a system event that has no consequence on the running health of the system. For example, user login notifications.

Note – The message categories are cumulative. For example, if you specify that you want to be notified about `notice` messages, you will also receive notification of `error` and `warning` messages. If you specify that you want to be notified about `information` messages, you will receive messages from all categories.



Caution – Use tab spaces to separate field entries when editing the `/etc/syslog.conf` file. If you use spaces rather than tab spaces to separate field entries, the array will not recognize your edits.

3. On the management host, use a text editor to edit the `syslog.conf` file in the working directory.

Separate entries with tab spaces. Edits are highlighted in the following example:

```
# syslog.conf
# facility.level action
# messages to local syslog file
*.notice /syslog

# messages to syslogd on another host Remote system logging
*.warn @remote-host
*.warn@129.234.56.73

# messages sent as SNMP traps
*.warn | snmp_trap 129.146.81.201
```

Note – Use the IP address in the `syslog.conf` file. If you want to add a host name, you must add a corresponding entry to the array's `/etc/hosts` file as described in the following section. The entry in the array's `/etc/hosts` file is used only for `syslog` control features on the array. Local utilities such as `ping` do not reference the array's `/etc/hosts` file. Therefore, IP addresses must be used with these utilities.

▼ To Edit the Array `/etc/hosts` File

You must edit the `/etc/hosts` file with the name of the management host and its IP address.

● **On the management host, use a text editor to edit the `/etc/hosts` file in the working directory.**

Separate the IP address from the name with a tab space. Edits are highlighted in the following example:

```
#hosts
#ip-address name
129.146.81.201 host-name
```

▼ To Transfer Files Back to the Array

After editing the `/etc/syslog.conf` and `/etc/hosts` files, transfer the files from the management host back to the array.

1. Start an `ftp` session from the management host to the array.

For example:

```
mngt_host:/:<15>ftp 129.146.81.201
Connected to 129.146.81.201.
220 chon-ji FTP server (SunOS 5.7) ready.
Name (129.146.81.201:root):
```

2. Log in to the array by typing `root` and your password at the prompts.

```
Name (129.146.81.201:root): root
331 Password required for root.
Password: password
230 User root logged in.
ftp>
```

3. Move to the `/etc` directory on the array.

```
ftp> cd /etc
250 CWD command successful.
ftp>
```

4. Type `binary` to set the transfer mode.

5. Copy the edited `syslog.conf` file from your working directory to the `/etc` directory on the array.

```
ftp> put syslog.conf
200 PORT command successful.
150 Binary data connection for syslog.conf (129.150.47.101,1031) (162 bytes).
226 Binary Transfer complete.
162 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

6. Copy the edited `hosts` file from your working directory to the `/etc` directory on the array.

```
ftp> put hosts
200 PORT command successful.
150 Binary data connection for hosts (129.150.47.101,1034) (47 bytes).
226 Binary Transfer complete.
47 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

7. Exit the ftp session.

```
ftp>
quit
221 Goodbye.
mngt_host::
```

8. On the array, start message traffic to the host.

The `syslogd` must be redirected at the array.

```
:::<23>set logto *
```

▼ To Edit the Management Host `/etc/syslog.conf` File

You must edit the `/etc/syslog.conf` file on the management host (the host used for monitoring and administering one or more arrays through the Ethernet connection) to include a log file for capturing error messages.

- On the management host, add a line with the name of the log file to the host's `/etc/syslog.conf` file, as shown in bold type in the following example:

Note – You must use a tab space between *log-filename* and `/var/adm/messages`.

```
#ident"@(##)syslog.conf1.496/10/11 SMI"/* SunOS 5.0 */
#
# Copyright (c) 1991-2001, by Sun Microsystems, Inc.
#
# syslog configuration file.
#
# This file is processed by m4 so be careful to quote (``) names
# that match m4 reserved words. Also, within ifdef's, arguments
# containing commas must be quoted.
#
*.err;kern.notice;auth.notice/dev/console
*.err;kern.debug;daemon.notice;mail.crit/var/adm/messages

*.alert;kern.err;daemon.erroperator
*.alert      root

*.emerg      *

# if a non-loghost machine chooses to have authentication messages
# sent to the loghost machine, un-comment out the following line:
#auth.noticeifdef('LOGHOST', /var/log/authlog, @loghost)

mail.debugifdef('LOGHOST', /var/log/syslog, @loghost)

#
# non-loghost machines will use the following lines to cause "user"
# log messages to be logged locally.
#
ifdef('LOGHOST', ,
user.err    /dev/console
user.err    /var/adm/messages
user.alert  'root, operator'
user.emerg  *
)
local7.notice /var/adm/messages.t300
```

In this example, all arrays that use this host for logging will have their messages dumped to `/var/adm/messages.t300`. The file name of the log file might be different depending on the type of array management software you are using.

Note – If you are using the StorTools™ product in the Solaris software environment to monitor the array, verify that you are using the correct `/var/adm/messages` file name. Refer to the StorTools documentation for the name of this message file.

Configuring the Array for SNMP Notification

The Sun StorEdge T3+ array can provide remote notification of array events to designated hosts using Simple Network Management Protocol (SNMP) traps. To enable SNMP notification, you must edit files on the array to configure system message logging. Because you cannot edit files on the array, you must use the `ftp` command to send them to a host to make the edits and then use the `ftp` command to return them to the array.

This procedure consists of the following tasks:

- “To Transfer the Array Files to the Management Host” on page 63
- “To Edit the Array `/etc/syslog.conf` File” on page 65
- “To Edit the Array `/etc/hosts` File” on page 67
- “To Transfer Files Back to the Array” on page 68

Note – For information on interpreting system messages, see Chapter 13.

▼ To Transfer the Array Files to the Management Host

1. Start an `ftp` session from the management host to the array.

For example:

```
mngt_host: /:<15>ftp 129.146.81.201
Connected to 129.146.81.201.
220 chon-ji FTP server (SunOS 5.7) ready.
Name (129.146.81.201:root):
```

2. Log in to the array by typing `root` and your password at the prompts.

```
Name (129.146.81.201:root): root
331 Password required for root.
Password: password
230 User root logged in.
ftp>
```

3. Move to your working directory on the management host.

For example:

```
ftp> lcd /tmp
Local directory now /tmp
ftp>
```

4. Move to the `/etc` directory on the array.

```
ftp> cd /etc
250 CWD command successful.
ftp>
```

5. Type `binary` to set the transfer mode.

6. Copy the `syslog.conf` file from the `/etc` directory on the array to your working directory.

```
ftp> get syslog.conf
200 PORT command successful.
150 Binary data connection for syslog.conf (129.150.47.101,1031) (162 bytes).
226 Binary Transfer complete.
162 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

7. Copy the `hosts` file from the `/etc` directory on the array to your working directory.

```
ftp> get hosts
200 PORT command successful.
150 Binary data connection for hosts (129.150.47.101,1034) (47 bytes).
226 Binary Transfer complete.
47 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

8. Exit the ftp session.

```
ftp>
quit
221 Goodbye.
mngt_host::
```

▼ To Edit the Array `/etc/syslog.conf` File

You must edit the `/etc/syslog.conf` file to include the category of messages that you want to have logged with the IP address of the management host that will be receiving these messages.

1. Determine the IP address and host name of the management host that will be logging messages.

Consult with your system administrator if you do not have this information.

2. Determine which category of messages you would like to receive from the array.

Messages generated by the array are grouped into four categories in the order of severity, as described in TABLE 7-2.

TABLE 7-2 Message Categories

Category	Description
Error	Indicates a critical system event requiring immediate user intervention or attention, such as the inability to flush the write-behind cache.
Warning	Indicates a serious system event requiring eventual user intervention. An example would be a disk drive being disabled.
Notice	Indicates a system event that may lead to a more serious condition in the future. An example would be the occurrence of hard errors corrected via parity replacement.
Information	Indicates a system event that has no consequence on the running health of the system. An example would be user login notifications.

Note – The message categories are cumulative. For example, if you specify that you want to be notified about `notice` messages, you will also receive notification of `error` and `warning` messages. If you specify that you want to be notified about `information` messages, you will receive messages from all categories.



Caution – Use tab spaces to separate field entries when editing the `/etc/syslog.conf` file. If you use spaces rather than tab spaces to separate field entries, the array will not recognize your edits.

3. On the management host, use a text editor to edit the `syslog.conf` file in the working directory.

Separate entries with tab spaces. Edits are highlighted in the following example:

```
# syslog.conf
# facility.level action
# messages to local syslog file
*.notice /syslog

# messages to syslogd on another host
*.warn @remote-host
*.warn@129.234.56.73

# messages sent as SNMP traps
*.warn | snmp_trap 129.146.81.201
```

SNMP notification

Note – You can use either the IP address or a host name in the `syslog.conf` file. If you use a host name, a corresponding entry must be present on the array's `/etc/hosts` file, as described in the next section.

▼ To Edit the Array `/etc/hosts` File

You must edit the `/etc/hosts` file with the name of the management host and its IP address.

- On the management host, use a text editor to edit the `/etc/hosts` file in the working directory.

Separate the IP address from the name with a tab space.

```
#hosts
#ip-address name
129.146.81.201 host-name
```

IP address

▼ To Transfer Files Back to the Array

After editing the `/etc/syslog.conf` and `/etc/hosts` files, transfer the files from the management host back to the array.

1. Start an `ftp` session from the management host to the array.

For example:

```
mngt_host: /: <15> ftp 129.146.81.201
Connected to 129.146.81.201.
220 chon-ji FTP server (SunOS 5.7) ready.
Name (129.146.81.201:root):
```

2. Log in to the array by typing `root` and your password at the prompts.

```
Name (129.146.81.201:root): root
331 Password required for root.
Password: password
230 User root logged in.
ftp>
```

3. Move to the `/etc` directory on the array.

```
ftp> cd /etc
250 CWD command successful.
ftp>
```

4. Copy the edited `syslog.conf` file from your working directory to the `/etc` directory on the array.

```
ftp> put syslog.conf
200 PORT command successful.
150 Binary data connection for syslog.conf (129.150.47.101,1031) (162 bytes).
226 Binary Transfer complete.
162 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

5. Copy the edited `hosts` file from your working directory to the `/etc` directory on the array.

```
ftp> put hosts
200 PORT command successful.
150 Binary data connection for hosts (129.150.47.101,1034) (47 bytes).
226 Binary Transfer complete.
47 bytes received in 1 seconds (0 Kbytes/s)
ftp>
```

6. Exit the ftp session.

```
ftp>
quit
221 Goodbye.
mngt_host::
```

7. On the array, start message traffic to the host.

The `syslogd` must be redirected at the array.

```
:::<23>set logto *
```


Volume Slicing, Mapping, and Masking Configuration

Firmware version 2.1 allows for multiple volume slicing (more than one slice per volume) which is similar to partitioning a hard drive of a desktop computer. This new feature allows greater flexibility to address storage in smaller increments, take advantage of larger capacity drives, and provides ease of scaling to suit individual applications.

Topics in this chapter include:

- “Volume Slicing Limitations” on page 73
- “Specifying Volume Slicing” on page 75
- “Specifying LUN Mapping” on page 89
- “Specifying LUN Masking” on page 92

Volume Configuration Limitations

A Sun StorEdge T3+ array consists of nine disk drives that, as discussed in Chapter 2, can be configured into one or two volumes. These volumes are presented to the host by the Sun StorEdge T3 array controller 2.1 software as a logical unit number (LUN). Thus, the application host does not see individual disk drives.

The following rules govern volume creation and reconfiguration.

1. Volumes must consist of whole drives.
2. A drive cannot be partitioned into different volumes.
3. A maximum of two volumes are allowed per T3 array unit.
4. A volume must equate to one drive group. A drive group is the exact equivalent of one RAID level. A RAID level is either striped (RAID 0 or RAID 5) or striped and mirrored (RAID 1+0). Thus, a drive group using RAID 5 can consist of either 3, 4, 5, 6, 7, 8, or 9 drives, and a RAID 1+0 drive group must consist of at least 2 drives. See *Sun STorEdge T3+ Array Installation and Configuration Manual* for an explanation of the RAID levels.
5. If drive nine is used as a hot-spare, it must be specified when a given volume is created. If you create two volumes, they can both use drive nine as a hot-spare if the hot-spare is designated when they are created. If you want to designate a hot-spare with an existing volume, that volume must be deleted and then recreated with a hot-spare (drive nine) designated.

Before creating or reconfiguring volumes in a T3 array, be sure to read “Volume Slicing Limitations” on page 73 and “T3 Array Configuration Considerations” on page 73. To create a volume, see “To Create a Logical Volume” on page 39.

Volume Slicing Limitations

1. The volume slicing commands apply only the controller firmware version 2.1 and later.
2. Volume slicing must be enabled with the `sys enable_volslice` command. See “To Enable Volume Slicing, Mapping, and Masking Operations” on page 76.
3. A maximum of 16 slices are allowed per T3 array workgroup or enterprise configuration.
4. The minimum slice size is 1 Gbyte and must be rounded on a RAID stripe boundary. If a RAID stripe is 9 blocks, and 1 Gbyte were to equal 100 blocks, then the minimum slice size would be $99 + 9$ blocks, or 108 blocks.
5. All slice sizes must be in multiples of the minimum slice size.
6. The maximum slice size is dependent on the volume size. All space in a volume need not be used.

T3 Array Configuration Considerations

When you are configuring a T3 array, consider these four parameters:

1. Will drive nine be used as a hot-spare? For a hot-spare discussion, see *Sun StorEdge T3+ Array Installation and Configuration Manual*.
2. How many volumes do you need—one or two? How many drives will be in each volume? Choosing between one or two volumes per T3 array requires balancing considerations of capacity, performance, and availability.

A configuration of a single volume per T3 array provides better performance than dual volumes because two volumes require additional administrative management overhead. However, an availability trade-off is associated with higher performance. For example, a one-volume T3 array might have a RAID 5 stripe of eight drives, so a failure of one of the eight drives results in a longer reconstruction time than if one drive within one of two four-drive RAID 5 volumes were to fail. This also leaves the eight-drive T3 array in a vulnerable state because a subsequent failure during the reconstruction of another drive within the same stripe results in a loss of data.

3. What RAID level is desired for each volume? See TABLE 8-1 and the *Sun StorEdge T3+ Array Installation and Configuration Manual* for a RAID discussion.

TABLE 8-1 T3+ Array Volume and Drive Configurations Available per RAID Level

RAID Level	Volume and Drive Configurations
0	<p>Single or two volumes—Nine drives maximum or eight plus one hot-spare; two drives minimum</p> <p>RAID 0 is advised only where performance is the sole qualifier or in conjunction with an external form of data protection such as using a host-based mirroring scheme across two Sun StorEdge T3+ arrays.</p>
1+0	<p>Single or two volumes—Nine drives maximum or eight plus one hot-spare; two drives minimum</p>
5	<ul style="list-style-type: none"> • Single volume—Nine drives maximum or eight plus one hot-spare; three drives minimum • Two volumes—Any combination of drives on the two volumes as long as at least three drives (exclusive of a hot-spare) are used in any one volume using RAID 5

4. How many total slices are desired, and how are they to be mapped to volumes? Consider that a failure within a volume affects all slices of that volume.

Specifying Volume Slicing

Firmware version 2.1 allows for volume slicing. This means that a volume can be partitioned into variable size slices and in effect allow multiple LUNs per workgroup or enterprise configuration. For more information on the individual commands, see Chapter 12.

This section contains the following topics:

- “To Enable Volume Slicing, Mapping, and Masking Operations” on page 76
- “To Reconfigure a Volume by Adding Slices After Volume Slicing Has Been Enabled” on page 79
- “To Create a Volume Slice” on page 88
- “To Remove a Volume Slice” on page 88

▼ To Enable Volume Slicing, Mapping, and Masking Operations

Sun StorEdge T3+ arrays might have one or two volumes set up. When you enable volume slicing, a default slice for each existing volume is created having the size of that volume. The resulting default slices are mapped to LUNs as 1:1 (one slice mapped to one LUN). The resulting default permissions are read/write.

Slicing a slice is not allowed. Therefore, if you want more than two slices, you must first remove the one or two default slices. If this is done, you must first backup any existing data before removing the slices since this will make the data inaccessible. The following procedure explains how to set up the Sun StorEdge T3+ array to do this and allow for volume slicing, mapping, and masking operations.

- Use the `sys enable_volslice` command on the master controller unit to enable volume slicing.

```
#!/usr/sbin/sys enable_volslice
WARNING - Once enabled this feature cannot be disabled
continue? Y
Initializing volslice configuration...
Volume slicing is enabled now!
2 out of Max. 16 slices created, 14 available. (example)
```



Caution – A procedure exists to disable the volume slicing feature; be aware that:

- The process is involved; detailed administrator knowledge is required.
- The risk of data loss exists.

If you still want to revert to the prior state before enabling volume slicing, perform one of the following two procedures.



Caution – If the default slices are removed and new slices are created (using the procedure described in “To Reconfigure a Volume by Adding Slices After Volume Slicing Has Been Enabled” on page 79) the data will be lost when you reset firmware version 2.0.1 or 2.1 to the state before enabling volume slicing.

▼ To Reset the Sun StorEdge T3+ Array Firmware Version 2.1

1. Back up all data onto a safe location.



Caution – The `boot -c` command destroys all system configuration information and resets the configuration to the factory default. Record this information before issuing the `boot -c` command.

2. Use the `boot -c` command to clear the system information and reboot.

```
:/: boot -c
Clear disk configuration, then system will reset, are you sure?
[N]: Y
```

The system automatically reboots.

3. Restore the prior system configuration.
4. Re-create the previous one or two volumes.
5. Restore the data.

▼ To Reset the Sun StorEdge T3+ Array Firmware Version 2.0.1

1. Back up all data onto a safe location.
2. Use the `volslice remove` command to remove all slices from the one or two volumes.

```
:/: volslice remove -v volume_name
WARNING - All slices in the given volume will be deleted.
Continue? [N]: Y
```

3. Use the `boot -i` command to download the firmware.

```
:/: boot -i firmware_binary
```

Where *firmware_binary* is the prior firmware file. For example, `nb201.bin`.

4. Use the `reset` command to reboot the Sun StorEdge T3+ array by typing:

```
:/: reset -y
```

5. Use the `boot -c` command to clear the system information and reboot.



Caution – The `boot -c` command destroys all system information.

```
:/: boot -c
Clear disk configuration, then system will reset, are you sure?
[N]: Y
```

The system automatically reboots.

6. Restore the prior one or two backed-up volumes.
7. Restore the data.

▼ To Reconfigure a Volume by Adding Slices After Volume Slicing Has Been Enabled

This procedure includes the following assumptions:

- Each T3+ array of the enterprise configuration has two volumes for a total of four volumes.
- The controller firmware has been upgraded to version 2.1.
- Volume slicing has been enabled with the `sys enable_volslice` command issued to the master controller unit. This results in slices `s0`, `s1`, `s2`, and `s3` for the four volumes (`v0`, `v1`, `v2`, and `v3`).
- Three new slices will be created on volume 0. This first requires the removal of the default slice `s0` created by the `sys enable_volslice` command.
- The application host has two HBAs connected, as shown in FIGURE 8-1.
- The Sun StorEdge T3+ Array enterprise configuration are connected via a private loop configuration.
- Sun StorEdge Traffic Manager Software is not installed or is disabled.

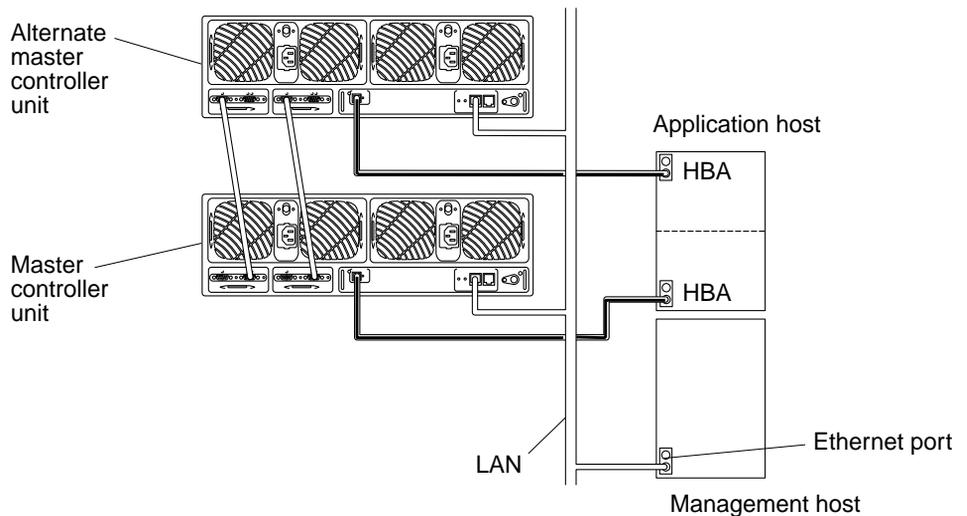


FIGURE 8-1 Sun StorEdge T3+ Array Enterprise Configuration Example

▼ To Remove the Default Slice and Create the Desired Slices

1. Back up all data on slice 0 (s0) onto a safe location(s).

Note – With the three slices to be created in this example, you must ensure that the available space on one of the newly created slices is sufficient to restore the backed-up data onto.

2. Use the `volslice remove` command to remove the default slice from volume 0.

```
:/: volslice remove s0
WARNING - The slice will be deleted.
Continue ? [N]: Y
```

Note – The default slices and names established by the `sys enable_volslice` command are `s0` for volume 0 and `s1` for volume 1.

3. Use the `volslice create` command to create the desired slices.

For example, to create a slice with a size of 5 Gbytes, a second slice with 15 Gbytes, and a third slice with 10 Gbytes on volume 0, type the following:

Note – This example shows what happens if you type a slice name that already exists—you are prompted to select another slice name.

```
:/: volslice create s1 -z 5GB v0
Slice name already exists
:/: volslice create s4 -z 5GB v0
4 out of Max. 16 slices created, 12 available.
:/:> volslice create s5 -z 15GB v0
5 out of Max. 16 slices created, 11 available.
:/: volslice create s6 -z 10GB v0
6 out of Max. 16 slices created, 10 available.
```

New slice names can be “Take1,” “slice1,” “Bob1,” or anything you desire. A slice name can be any 12 (maximum) alphanumeric characters except a blank space. The default permission of a newly created slice is set at *none*—the system will not permit any access by any host.

4. Use the `volslice list` command to verify the slice sizes.

```
:/: volslice list
Slice      Slice Num      Start Blk      Size Blks      Capacity      Volume
s1         0              0              104870400      5.014 GB      v1
-         -              104870400      190125327      68.3 GB       v1
s2         1              0              104870400      5.014 GB      v2
-         -              104870400      190125327      68.3 GB       v2
s3         2              0              104870400      5.014 GB      v3
-         -              104870400      190125327      68.3 GB       v3
s4         4              0              104870400      5.014 GB      v0
s5         5              104870400      120334782      15.020 GB     v0
s6         6              120334782      135335223      10.018 GB     v0
-         -              135335223      190125327      43.8 GB       v0
```

5. Use the `lun perm list` command to verify the permissions are set to none.

```
:/: lun perm list
Lun  Slice  WWN              Group Name      Group Perm      WWN Perm      Effective Perm
-----
4    4      default          --              --              none          none
5    5      default          --              --              none          none
5    6      default          --              --              none          none
1    1      default          --              --              none          none
2    2      default          --              --              none          none
3    3      default          --              --              none          none
-----
```

▼ To Obtain the Controller Card WWNs for each T3+ Array of the Enterprise Configuration

1. Use the `port list` command to display the WWN of each controller card of the T3+ array enterprise configuration.

```
:/: port list

port      targetid  addr_type  status  host  wwn
ulp1     5         hard      online  sun   50020f230000a74d
u2p1     6         hard      online  sun   50020f230000a8bf
```

2. Record the WWN of each T3+ array controller card.

3. Use the `lun default` command to assign read/write permission to all LUNs/slices.

This allows the host to detect all slices.

```
:/: lun default all_lun rw
```

4. Use the `lun perm list` command to verify all permissions are set at read/write.

```
:/: lun perm list
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
4	4	default	--	--	rw	rw
5	5	default	--	--	rw	rw
6	6	default	--	--	rw	rw
1	1	default	--	--	rw	rw
2	2	default	--	--	rw	rw
3	3	default	--	--	rw	rw

5. On the host, use the `format` command to display the paths to each LUN/slice.

```
# format
1. c3t5d0 <SUN-T300-0201 cyl 4680 alt 2 hd 7 sec 128>
   /ssm@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0/ssd@w50020f230000a74d,0
2. c3t5d1 <SUN-T300-0201 cyl 11703 alt 2 hd 7 sec 128>
   /ssm@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0/ssd@w50020f230000a74d,1
3. c3t5d2 <SUN-T300-0201 cyl 4680 alt 2 hd 7 sec 128>
   /ssm@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0/ssd@w50020f230000a74d,2
4. c4t6d0 <SUN-T300-0201 cyl 4680 alt L2 hd 7 sec 128>
   /ssm@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f230000a8bf,0
5. c4t6d1 <SUN-T300-0201 cyl 11703 alt 2 hd 7 sec 128>
   /ssm@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f230000a8bf,1
6. c4t6d2 <SUN-T300-0201 cyl 4680 alt 2 hd 7 sec 128>
   /ssm@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f230000a8bf,2
```

There are six paths—three LUNs/slices and two paths to each LUN/slice (through the two HBAs).

6. Confirm the two controller card WWNs for the two T3+ array units with the WWNs recorded from the output of Step 1 and 2 above.

These are listed in line 1 and 4 in the `format` command output:

- 50020f230000a74d
- 50020f230000a8bf

▼ To Obtain the WWN for the First Host HBA Port

1. On the host, use the `luxadm disp` command with the logical path of the first LUN/slice to obtain its physical path.

Lines “1.” through “3.” of the format command output show the logical device names for the ...a74d controller card. Use the first one with “s2” to indicate the entire LUN/slice (c3t5d0s2) and add it to complete the logical device name path (/dev/rdsk/c3t5d0s2).

```
# luxadm disp /dev/rdsk/c3t5d0s2
DEVICE PROPERTIES for disk: /dev/rdsk/c3t5d0s2
  Status (Port B):      O.K.
  Vendor:              SUN
  Product ID:          T300
  WWN (Node):          50020f230000a74d
  WWN (Port B):        50020f230000a8bf
  Revision:            0200
  Serial Num:          Unsupported
  Unformatted capacity: 136588.000 MBytes
  Write Cache:         Enabled
  Read Cache:          Enabled
  Minimum prefetch:    0x0
  Maximum prefetch:    0x0
  Device Type:         Disk device
  Path(s):
    /dev/rdsk/c3t5d0s2
    /devices/ssp@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0/ssd@w50020f230000a74d,0
```

The physical LUN/slice path in this example is:

```
/devices/ssp@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0
/ssd@w50020f230000a74d,0.
```

2. Replace the `/ssd@w50020f230000a74d,0` portion of the physical LUN/slice path with `:devctl` to obtain the physical HBA controller path:

```
/devices/ssp@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0:devctl
```

3. On the host, use the luxadm command with the -e dump_map option and the physical HBA controller path to display the WWN for the first HBA port.

Note – The luxadm -e dump_map command can only be used with Fibre Channel controller ports.

```
# luxadm -e dump_map
/devices/ssm@0,0/pci@1c,600000/pci@1/SUNW,qlc@5/fp@0,0:devctl
Pos  Port_ID  Hard_Addr  Port WWN          Node WWN          Type
0    1044d5   1000d5     50020f230000091f  50020f200000091f  0x0(Disk device)
1    1046d1   1000d1     50020f230000064a  50020f200000064a  0x0(Disk device)
2    1047cd   1000cd     50020f2300004251  50020f2000004251  0x0(Disk device)
3    104800   0          210100e08b208900  200100e08b247212  0x1f(Unknown
Type,Host Bus Adapter)
```

In this example, the first host HBA port is identified as 210100e08b208900. Record the WWN number for use in “To Set Up Access for the Desired Host HBA Ports” on page 86.

▼ To Obtain the WWN for the Second Host HBA Port

1. On the host, use the `luxadm disp` command with the logical path of the second LUN/slice to obtain its physical path.

Lines “4.” through “6.” of the format command output show the logical device names for the ...a8bf controller card. Use the first one with “s2” to indicate the entire LUN/slice (c4t6d0s2) and add it to complete the logical device name path (/dev/rdsk/c4t6d0s2).

```
# luxadm disp /dev/rdsk/c4t6d0s2
DEVICE PROPERTIES for disk: /dev/rdsk/c4t6d0s2
  Status(Port B):      O.K.
  Vendor:              SUN
  Product ID:         T300
  WWN(Node):          50020f230000a74d
  WWN(Port B):        50020f230000a8bf
  Revision:           0200
  Serial Num:         Unsupported
  Unformatted capacity: 136588.000 MBytes
  Write Cache:        Enabled
  Read Cache:         Enabled
  Minimum prefetch:   0x0
  Maximum prefetch:   0x0
  Device Type:        Disk device
  Path(s):
  /dev/rdsk/c4t6d0s2
  /devices/ssp@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0/ssd@w50020f230000a8bf,0
```

The physical LUN/slice path in this example is:

```
/devices/ssp@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0
/ssd@w50020f230000a8bf,0.
```

2. Replace the `/ssd@w50020f230000a8bf,0` portion of the physical LUN/slice path with `:devctl` to obtain the physical HBA controller path:

```
/devices/ssp@0,0/pci@1d,700000/pci@1/SUNW,qlc@4/fp@0,0:devctl
```

3. On the host, use the `luxadm` command with the `-e dump_map` option and the physical HBA controller path to display the WWN for the first HBA port.

Note – The `luxadm -e dump_map` command can only be used with Fibre Channel controller ports.

```
# luxadm -e dump_map
/devices/ssm@0,0/pci@1d,70000/pci@1/SUNW,qlc@4/fp@0,0:devctl
Pos  Port_ID  Hard_Addr  Port WWN          Node WWN          Type
0    1044d5   1000d5     50020e230000091b  50020e200000091b  0x0(Disk device)
1    1046d1   1000d1     50020e230000064c  50020e200000064c  0x0(Disk device)
2    1047cd   1000cd     50020e2300004251  50020e2000004251  0x0(Disk device)
3    104800   0          210000e08b04aa9a  200100e08b247212  0x1f(Unknown
Type,Host Bus Adapter)
```

In this example, the second HBA port is identified as 210000e08b04aa9a. Record the WWN number for use in “To Set Up Access for the Desired Host HBA Ports” on page 86.

▼ To Set Up Access for the Desired Host HBA Ports

1. On the T3+, use the `lun default` command to remove the read/write permission from all slices.

This allows setting a specific permission for an HBA port for LUN masking.

```
:/: lun default all_lun none
```

2. Use the `lun perm` command to:

- Set the LUN 0 permission as desired (read/write in the example below).
- Set up access to the desired application host HBA port WWN with the named permission.

```
:/: lun perm lun 0 rw wwn 210100e08b208900 (first HBA port WWN)
```

3. Verify the result with the `lun perm list` command.

```
:/: lun perm list
Lun  Slice  WWN                Group Name  Group Perm  WWN Perm  Effective Perm
-----
0    4      210100e08b208900  --         --         rw       rw
0    5      default           --         --         none     none
0    6      default           --         --         none     none
1    1      default           --         --         none     none
2    2      default           --         --         none     none
3    3      default           --         --         none     none
-----
```

4. Repeat Steps 2 and 3 for each LUN.

When you are finished with all LUNs, restore the data onto the slices.

▼ To Create a Volume Slice

1. Use the `volslice create` command to create slice 1 with a size of 5 Gbytes on volume 2.

Note – If you issue the `volslice create` command without any parameters, you will be interactively prompted to specify the parameters.

```
:/: volslice create s1 -z 5GB v2
1 out of Max. 16 slices created, 15 available.
```

2. Use the `volslice list` command to list the details of volume 2 to verify the slice.

```
:/: volslice list -v v2
Slice      Slice Num      Start blk      Size blks      volume
-          -              0              1023           v2
s1         0              1024           4966680        v2
-          -              4966680        87418240       v2
```

▼ To Remove a Volume Slice

- Use the `volslice remove` command to either remove a specified slice or remove all slices of a specified volume.

```
:/: volslice remove s4
WARNING - The slice will be deleted.
Continue ? [N]: Y

:/: volslice remove -v v1
WARNING - All slices in the given volume will be deleted.
Continue ? [N]: Y
```

Specifying LUN Mapping

The LUN mapping commands apply only to the controller firmware version 2.1 and later with volume slicing enabled. For more information on the individual commands, see Chapter 12.

This section contains the following topics:

- “To Display the LUN Map” on page 89
- “To Add an Entry to the LUN Map” on page 90
- “To Remove an Entry From the LUN Map” on page 91

▼ To Display the LUN Map

1. Use the `lun map list` command to display the LUN map.

```
:/: lun map list
Lun No   Slice No
0        0
1        1
2        2
3        3
```

▼ To Add an Entry to the LUN Map

1. Use the `lun map add` command to map LUN 4 to slice 3.

```
:/: lun map add lun 4 slice 3
```

2. Use the `lun map list` command to verify the result.

```
:/: lun map list
Lun No    Slice No
0         0
1         1
2         2
3         4
4         3
```

▼ To Remove an Entry From the LUN Map

1. Use the `lun map rm` command to remove LUN mapping from LUN 4 to slice 3.

```
:/: lun map rm lun 4 slice 3
```

2. Use the `lun map list` command to verify the result.

```
:/: lun map list
Lun No    Slice No
0         0
1         1
2         2
3         4
```

Specifying LUN Masking

The LUN masking commands apply only to the controller firmware version 2.1 and later with volume slicing enabled.

LUN masking is a key component of building large storage systems. It eases storage administration while allowing for a more secure environment. When a LUN is masked from a host, it is not available to be configured and thus prevents the LUN from being assigned to multiple hosts accidentally. This becomes critical for large network configurations. In such networks, if each host could access all of the LUNs, it would be difficult to administratively manage the host permissions.

Each host bus adapter (HBA) port is assigned a unique world-wide name (WWN) by the HBA manufacturer. A WWN is composed of 16 alphanumeric characters. After the HBA is installed on a host, the host system software (for example, the `luxadm` administration program in the Sun Enterprise Network Array in the Solaris operating environment) can be used to identify the WWN for each HBA port.

The Sun StorEdge T3+ array controller firmware version 2.1 does not recognize (*register*) a WWN until the WWN is recorded by the `lun perm` or `hwnn add` commands.

WWNs can be assigned to a specific LUN or LUNs. WWNs can also be grouped and assigned to a specific LUN or LUNs. This is referred to as *LUN masking*. LUN masking requires the verification of each host I/O by checking the host WWN against the LUN permissions. This allows for configuration planning for increased security.

For more information on the individual commands, see Chapter 12.

This section contains the following topics:

- “To Find the Default LUN Permissions and System-Assigned WWNs” on page 93
- “To Set a Specific LUN Permission” on page 94
- “To View All Registered WWNs” on page 94
- “To Change the Default LUN Permissions” on page 96
- “To Define a Host WWN Group” on page 97
- “To Set the Permissions of a LUN for All Members of a Group” on page 98
- “To Remove the Registered Status of a WWN” on page 100
- “To Remove a WWN From a Specific Group” on page 101
- “To Remove All WWNs From a Specific Group” on page 102

▼ To Find the Default LUN Permissions and System-Assigned WWNs

- Use the `lun perm list` command to find the default LUN permissions.

For each LUN, the default access permissions are followed by any exceptions.

```
:/: lun perm list
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	rw	rw
0	0	20020678ff345678	G1	none	ro	ro
1	1	default	--	--	ro	ro
1	1	20020678ff345678	G1	ro	rw	rw
1	1	20020678ee345678	G1	ro	none	ro
1	1	20020678ee345678	G2	ro	none	ro
1	1	20020678ab345678	G2	ro	none	ro

▼ To Set a Specific LUN Permission

1. Use the `lun perm` command to set up LUN 2 with read/write access to WWN 20020da445678901.

```
:/: lun perm lun 2 rw wwn 20020da445678901
```

Note – The `lun perm` command “registers” the specified WWN, if it is not already registered.

2. Use the `lun perm list` command to verify the result.

```
:/: lun perm list wwn 20020da445678901
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	none	none
0	0	20020da445678901	--	--	rw	rw
1	1	default	--	--	none	none
1	1	20020da445678901	--	--	rw	rw
2	2	default	--	--	none	none
2	2	20020da445678901	--	--	rw	rw

▼ To View All Registered WWNs

1. Use the `lun wwn list` command.

```
:/: lun wwn list
List of all registered WWNs
-----
20020da445678901
20020da445678901
```

Setting LUN Permission with Host Multipathing

The Sun StorEdge T3+ arrays are redundant devices that automatically reconfigure whenever a failure occurs on a T3+ array internal component. To allow for the automatic reconfiguration to take effect, a host-based software solution—for example, Sun StorEdge Traffic Manager Software (STMS) or VERITAS Dynamic Multipathing (DMP) and a redundant data path are also required. This redundant path is referred to as *host multipathing*. See FIGURE 8-2.

During normal operation, data I/O takes place on the host channel connected to the controller that owns the LUNs. This path is a *primary path*. During fail-over operation, the multipathing software directs all I/Os to the alternate channel's controller. This path is the *fail-over path*.

When a controller card in the master controller unit fails, the alternate master controller unit becomes the master. When the failed controller is replaced, the new controller immediately boots, goes online, and the original master controller unit becomes the alternate master controller unit. The former alternate master controller unit remains the master controller unit.

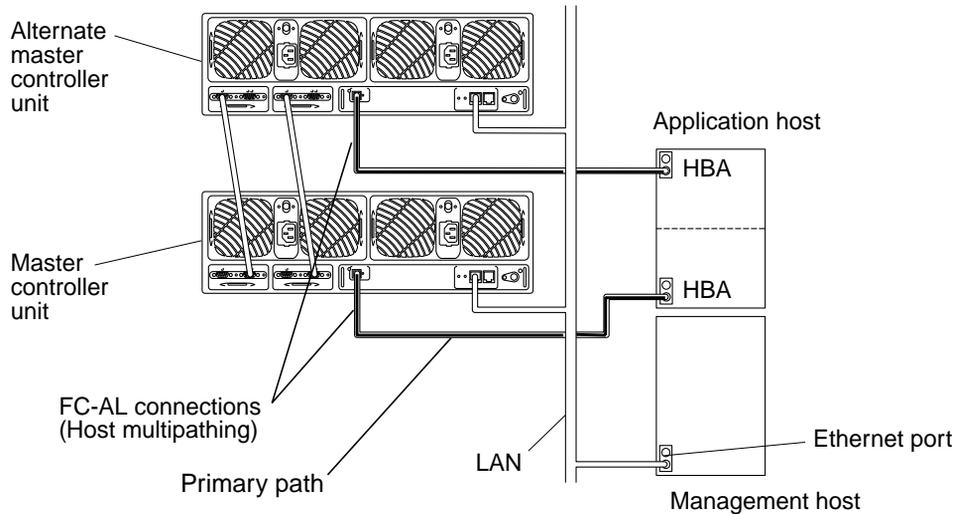


FIGURE 8-2 Host Multipathing Configuration

When a system is configured to incorporate both LUN masking and host multipathing, you must ensure that all WWNs of all HBAs (or, if applicable, both ports being used on a single HBA) are set with the same permissions to permit the desired fail-over communications. The easiest way to do this is to group the WWNs

with a common group name and then set the permission for that group. See “To Define a Host WWN Group” on page 97 and “To Set the Permissions of a LUN for All Members of a Group” on page 98.

▼ To Change the Default LUN Permissions

1. Use the `lun default` command to establish the default permission to LUN 3 as read only.

```
:/: lun default lun 3 ro
```

Note – The `lun default` command changes the permissions of only nonregistered WWNs.

2. Use the `lun perm list` command to verify the result.

```
:/: lun perm list lun 3
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
3	3	default	--	--	none	none
3	3	20020678ff345678	--	--	ro	ro

▼ To Define a Host WWN Group

If all HBAs on a host will share the same access permissions, you can use this procedure to define a host WWN group. The two methods are shown merely as an example.

Note – If you issue the `hwwn add` command without any parameters, you will be interactively prompted to specify the parameters.

1. Use the `hwwn add` command to add the WWN 20020678ef345678 to group “node1”

■ Interactively

```
:/: hwwn add
Please enter Group Name: node1
Please enter WWN: 20020678ef345678
Add more WWN? [n]: N
```

■ Manually

```
:/: hwwn add node1 wwn 20020321ff345678
```

Note – The `hwwn add` command “registers” the specified WWN, if it is not already registered. The permission of the existing group also becomes the permission of the added WWN.

2. Use the `lun perm list` command to verify the result.

```
:/: lun perm list grp node1
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	rw	rw
0	0	20020678ff345678	node1	none	ro	ro
1	1	default	--	--	ro	ro
1	1	20020678ff345678	node1	ro	rw	rw
1	1	20020678ff345678	node1	ro	none	rw
2	2	default	--	--	rw	none

▼ To Set the Permissions of a LUN for All Members of a Group

With a host WWN defined for a group, this procedure can be used to set permissions using the host WWN group name.

1. Use the `lun perm list` command to display the current LUN permissions.

```
:/: lun perm list
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	rw	rw
0	0	20020678ff345678	G1	none	ro	ro
1	1	default	--	--	ro	ro
1	1	20020678ff345678	G1	ro	rw	rw
1	1	20020678ee345678	G1	ro	none	ro
1	1	20020678ee345678	node2	rw	none	rw
1	1	20020678ab345678	node2	rw	none	rw

2. Use the `lun perm` command to set read-only permissions on LUN 1 for group "node2."

```
:/: lun perm lun1 ro grp node2
```

3. Use the `lun perm list` command to verify the result.

```
:/: lun perm list
Lun  Slice  WWN                Group Name  Group Perm  WWN Perm  Effective Perm
-----
0    0      default            --          --          rw        rw
0    0      20020678ff345678  G1         none        ro        ro
1    1      default            --          --          ro        ro
1    1      20020678ff345678  G1         ro          rw        rw
1    1      20020678ee345678  G1         ro          none       ro
1    1      20020678ee345678  node2      ro          none       ro
1    1      20020678ab345678  node2      ro          none       ro
-----
```

▼ To Remove the Registered Status of a WWN

1. Use the `lun wwn list` command to display the registered WWNs.

```
:/: lun wwn list
List of all registered WWNs
-----
2002067890123456
2002078901234567
2002089012345678
2002090123456789
```

2. Use the `lun wwn rm` command to remove the `2002067890123456` WWN.

```
:/: lun wwn rm wwn 2002067890123456
2002067890123456 will get default access right.
Do you want to remove? (Y/N)[n]:Y
```

3. Use the `lun wwn list` command to verify the result.

```
:/: lun wwn list
List of all registered WWNs
-----
2002078901234567
2002089012345678
2002090123456789
```

▼ To Remove a WWN From a Specific Group

1. Use the `hwwn list` command to display all the WWN(s) in a specified group.

```
:/: hwwn list node1
WWN in this group node1
-----
20020678ef345678
20020321ff345678
-----
** Total 2 entries **
```

2. Use the `hwwn rm` command to remove both WWNs from the group “node1.”

```
:/: hwwn rm node1 wwn 20020678ef345678
:/: hwwn rm node1 wwn 20020321ff345678
```

Note – If you issue the `hwwn rm` command without any parameters, you will be interactively prompted to specify the parameters.

▼ To Remove All WWNs From a Specific Group

1. Use the `hwn listgrp` command to display all registered group names.

```
:/:<#> hwn listgrp
Group Name
-----
node1
node2
node3
-----
** Total 3 entries **
```

2. Use the `hwn list` command to display all WWN(s) in a specified group.

```
:/:<#> hwn list node1
WWN in this group node1
-----
20020678ef345678
20020321ff345678
-----
** Total 2 entries **
```

3. Use the `hwn rmgrp` command to remove all WWN(s) from a specified group.

```
:/:<#> hwn rmgrp node1
```

4. Use the `hwwn list` command to verify the results.

```
:/:<#> hwwn list node1
WWN in this group node1
-----
-----
** Total 0 entries **
```


Network Time Protocol Configuration

These commands apply only to the controller firmware version 2.1 and later.

The Network Time Protocol (NTP) support feature allows a network system of Sun StorEdge T3+ arrays to automatically synchronize their internal system time with a designated NTP master Sun StorEdge T3+ array.

This chapter discusses the following topic:

- “Specifying Network Time Protocol” on page 106

Specifying Network Time Protocol

All commands are done through the T3+ array. For more information on the individual commands, see Chapter 12.

▼ To Set Up the NTP Server

1. Use the `set timezone` command to establish the time zone.

```
:/:<#> set timezone :America/Los_Angeles
```

2. Use the `set` command to confirm the time zone region setting.

```
:/:<7> set
bootmode      tftp
bootdelay     3
sn            103676
ip            10.4.31.28
netmask       255.255.255.0
gateway       10.4.31.1
tftpghost     10.4.31.83
tftpfile      t4ssl.bin
hostname      InUseByAlbert
timezone      :America/Los_Angeles
vendor        0301
model         501-5710-02(51)
revision      0200
logto         /Apr02
loglevel      3
rarp          on
mac           00:20:f2:01:94:fc
```

3. Use the `ntp server` command to set the IP address of the server that is running the NTP daemon.

```
:/:<#> ntp server 22.3.36.42
```

4. Use the `ntp interval` command to set up the polling interval in minutes of an NTP client to the NTP server.

```
:/:<#> ntp interval 5
```

In this example, the array polls the NTP server (updates the synchronization) every 5 minutes.

5. Use the `ntp` command to verify that the server and polling interval are set correctly.

```
:/:<#> ntp
server          22.3.36.42
poll            off
interval        5
```

The example above shows the NTP server's IP address, that polling is turned off, and the polling interval is set at 5 minutes.

6. Use the `ntp poll` command to turn on the NTP polling, if required.

```
:/:<#> ntp poll unicast
```

7. Use the `ntp stats` command to verify the NTP client (Sun StorEdge T3+ array) status.

```
:/:<#> ntp stats
lastpoll      Tue Feb 19 21:07:32 GMT 2002
server        22.3.36.42
offset        + 0.00413176
status        Successfully adjusted the time.
```

The example above shows the time of the last poll, the NTP server's IP address, the adjustment made the T3+ array internal clock, and the clock set success statement.

Fibre Channel Topology Configuration

The commands in this chapter apply only to the controller firmware version 2.1 and later.

Prior versions of controller firmware only supported a Fibre Channel loop topology with the Sun StorEdge T3+ array connectivity. With the 2.1 firmware release, a point-to-point topology is now also supported. This means that a Sun StorEdge T3+ array can now connect directly to a switch, hub, or host, with full fabric connectivity.

This chapter discusses the following topic:

- “Specifying Fibre Channel Topology” on page 110

Specifying Fibre Channel Topology

For more information on the individual commands, see Chapter 12. The section contains the following topics:

- “To Determine the Topology Firmware Mode” on page 110
- “To Change the Topology Firmware Mode” on page 111

▼ To Determine the Topology Firmware Mode

- Use the `sys list` command to determine the topology firmware mode setting of the T3+ array.

```
:/:<#> sys list
blocksize           : 64k
cache               : auto
mirror              : auto
mp_support          : none
naca                 : off
rd_ahead            : on
recon_rate          : med
sys memsize         : 128 MBytes
cache memsize       : 1024 MBytes
enable_volslice     : on
fc_topology         : auto
enable_volslice     : on
fc_topology         : auto
```

In the above example, the T3+ array Fibre Channel topology mode is set to automatic sense and set. This is the default setting. See TABLE 10-1 for an explanation of the various mode settings. For a full explanation of the `fc_topology` command, see “CLI Commands” on page 131.

TABLE 10-1 fc_topology mode settings

Parameter	Function
auto	Enables the controller firmware to automatically determine and set the topology firmware mode.
loop	Sets the loop topology firmware mode.
fabric_p2p	Sets the point-to-point topology firmware mode.

▼ To Change the Topology Firmware Mode

You should not need to perform this task except for troubleshooting purposes.

1. Set the topology firmware driver of the T3+ array as desired. For example, to set the driver to point-to-point, type:

```
:/:<#> sys fc_topology fabric_p2p
```

2. Wait at least 10 seconds before querying the current negotiated topology.
3. Use the `sys fc_topology` command to verify the current negotiated topology.

```
:/:<#> sys fc_topology  
Host port ulp1 Current Negotiated Topology=Fabric  
Point to Point, Port ID=100300
```


Advanced Array Troubleshooting

This chapter describes using the CLI to troubleshoot the array and its components.

This chapter contains the following sections:

- “To Identify FRU Failures” on page 114
- “Replacing FRUs” on page 119
- “To Identify Miscabled Partner Groups” on page 126
- “Identifying Data Channel Failures” on page 130

▼ To Identify FRU Failures

The first step in remedying a FRU failure is to determine which FRU has failed. You can do this by examining the output of the `fru stat` command and the messages in the `syslog` file. After evaluating the data, you can confirm the failing FRU and replace it.

1. Display a list of all installed FRUs and check their current status.

```

:/: fru stat

```

CTLR	STATUS	STATE	ROLE	PARTNER	TEMP			
ulctr	ready	enabled	master	u2ctr	32.0			
u2ctr	ready	enabled	alt master	ulctr	29.0			
DISK	STATUS	STATE	ROLE	PORT1	PORT2	TEMP	VOLUME	
uld1	ready	enabled	data disk	ready	ready	34	v0	
uld2	ready	enabled	data disk	ready	ready	34	v0	
uld3	ready	enabled	data disk	ready	ready	31	v0	
uld4	ready	enabled	data disk	ready	ready	34	v0	
uld5	ready	enabled	data disk	ready	ready	32	v0	
uld6	ready	enabled	data disk	ready	ready	38	v0	
uld7	ready	enabled	data disk	ready	ready	33	v0	
uld8	ready	enabled	data disk	ready	ready	34	v0	
uld9	ready	enabled	data disk	ready	ready	36	v0	
u2d1	ready	enabled	data disk	ready	ready	30	v1	
u2d2	ready	enabled	data disk	ready	ready	31	v1	
u2d3	ready	enabled	data disk	ready	ready	34	v1	
u2d4	ready	enabled	data disk	ready	ready	34	v1	
u2d5	ready	enabled	data disk	ready	ready	32	v1	
u2d6	ready	enabled	data disk	ready	ready	37	v1	
u2d7	ready	enabled	data disk	ready	ready	38	v1	
u2d8	ready	enabled	data disk	ready	ready	42	v1	
u2d9	ready	enabled	data disk	ready	ready	34	v1	
LOOP	STATUS	STATE	MODE	CABLE1	CABLE2	TEMP		
u2l1	ready	enabled	master	installed	-	30.0		
u2l2	ready	enabled	slave	installed	-	30.5		
u1l1	ready	enabled	master	-	installed	27.0		
u1l2	ready	enabled	slave	-	installed	29.5		
POWER	STATUS	STATE	SOURCE	OUTPUT	BATTERY	TEMP	FAN1	FAN2
u1pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u1pcu2	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu2	ready	enabled	line	normal	normal	normal	normal	normal

2. Review the output of the fru stat list for any fault or abnormal status messages.

See TABLE 6-2 for a list of possible FRU status messages.

3. Display the `syslog` file.

Use the `tail` command to display the last 10 lines of the file, which contain the most recent log entries. (See Chapter 13 for a description of the `syslog` format.)

```
#!/: tail syslog
Oct 26 16:24:15 ROOT[2]: W: u2ctr Hardware Reset (1000) occurred
Oct 26 16:24:16 ROOT[2]: N: u2ctr Waiting for configuration data from master...
Oct 26 16:24:31 MNXT[2]: N: u2ctr Initializing host port u2p1 ISP2100...
firmware status = 3
Oct 26 16:24:31 MNXT[2]: N: u2ctr Host port u2p1 TARGET_ID = 0x2 (ALPA = 0xe4)
Oct 26 16:24:58 ROOT[2]: N: u2ctr Starting psh...
Oct 26 16:24:58 ISR1[2]: N: u2ctr ISP2100[2] Received LIP(f7,e4) async event
Oct 26 16:24:58 FCC0[2]: N: u2ctr Port event received on port 3, abort 0
Oct 26 16:25:08 FCC2[2]: N: u2ctr (ITL 0 1 0 TT 20 TID 7910 OP 0) Target in Unit
Attention
```

For more extensive investigation, display the complete `syslog` file. You can do this in two ways:

- Using the `cat` command (that is, `cat syslog`), which could overwrite the terminal window.
- Moving the `syslog` file to the host using an FTP connection. See “To Transfer Array Files to the Management Host” on page 56 for instructions.

4. Use a text editor to examine the `syslog` file.

Look for Error and Warning messages—these indicate a serious problem with the unit. Error message abbreviations are: E—errors, W—warnings, N—notices, I—informational. The following is a partial output of a `syslog` file:

```
Oct 26 17:01:30 LPCT[1]: W: u2pcu1: Switch off
Oct 26 17:01:31 LPCT[1]: W: u2pcu1: Off
Oct 26 17:01:33 LPCT[1]: W: u2pcu1: DC not OK
Oct 26 17:01:36 LPCT[1]: E: u2pcu1: Battery not present
Oct 26 17:01:39 LPCT[1]: E: u2pcu1: Not present
Oct 26 17:01:39 TMRT[1]: E: u2pcu1: Missing; system shutting down in 30 minutes
Oct 26 17:01:40 TMRT[1]: E: u2ctr: Multiple Fan Faults; system shutting down in
30 minutes
Oct 26 17:01:42 LPCT[2]: E: u2pcu1: Not present
Oct 26 17:01:45 LPCT[1]: E: u2pcu1: Over temperature
Oct 26 17:01:45 LPCT[1]: W: u2pcu1: Switch off
Oct 26 17:01:46 LPCT[1]: W: u2pcu1: Off
Oct 26 17:01:46 LPCT[1]: E: u2pcu1: Battery not present
Oct 26 17:01:48 LPCT[1]: N: u2pcu1: Battery not OK
Oct 26 17:04:16 LPCT[1]: E: u2d4: Not present
Oct 26 17:04:16 TMRT[1]: E: u2d4: Missing; system shutting down in 30 minutes
```

Note – The array has limited space for `syslog` files. When the `syslog` file exceeds the file-size threshold (1 Mbyte), this file is moved to a file called `syslog.OLD`, overwriting any previous `syslog.old` image.

If you are monitoring an error condition, transfer the `syslog` file to a host to save the output and prevent any pertinent data from being deleted.

5. Compare `syslog` entries with the output of the `fru stat` command to see which FRU has failed.

In the following example, a warning (W:) message is shown in array unit 2 for drive 7 (u2d7). The `fru stat` output indicates a problem with u2d7. Based on this information, you would determine whether drive 7 in unit 2 needs to be repaired or replaced.

```
Oct 26 17:13:38 FCC0[1]: N: ulctr (ITL 7D 1 0 TT 20 TID 77A8 OP 0) Target in
Unit Attention
Oct 26 17:13:37 FCC0[2]: N: u2ctr (ITL 7D 2 1 TT 20 TID 77A8 OP 0) Target in
Unit Attention
Oct 26 17:14:12 ISR1[2]: W: u2d7 SCSI Disk Error Occurred (path = 0x0)
Oct 26 17:14:12 ISR1[2]: W: Sense Key = 0x2, Asc = 0x4, Ascq = 0x2
Oct 26 17:14:12 ISR1[2]: W: Sense Data Description = Logical Unit Not Ready,
Initializing CMD Required
Oct 26 17:14:13 ISR1[1]: N: ulctr ISP2100[1] Received LIP(f8,e8) async event
Oct 26 17:14:12 ISR1[2]: N: u2d7 SVD_DONE: Command Error = 0x3
Oct 26 17:14:12 ISR1[2]: N: u2d7 sid 0 stype 815 disk error 3
Oct 26 17:14:12 ISR1[2]: W: u2d7 SCSI Disk Error Occurred (path = 0x0)
Oct 26 17:14:15 SVDT[1]: N: 19 fcal ports were detected on l2
Oct 26 17:14:16 LPCT[1]: N: u2d7: Bypassed on loop 1
Oct 26 17:14:12 ISR1[2]: W: Sense Key = 0x2, Asc = 0x4, Ascq = 0x2
Oct 26 17:14:12 ISR1[2]: W: Sense Data Description = Logical Unit Not Ready,
Initializing CMD Required
Oct 26 17:14:12 ISR1[2]: N: u2d7 SVD_DONE: Command Error = 0x3
Oct 26 17:14:17 LPCT[1]: N: u2d7: Bypassed on loop 2
Oct 26 17:14:12 ISR1[2]: N: u2d7 sid 0 stype 405 disk error 3
Oct 26 17:14:12 ISR1[2]: W: u2d7 SCSI Disk Error Occurred (path = 0x0)
Oct 26 17:14:12 ISR1[2]: W: Sense Key = 0x2, Asc = 0x4, Ascq = 0x2
Oct 26 17:14:12 ISR1[2]: W: Sense Data Description = Logical Unit Not Ready,
Initializing CMD Required
Oct 26 17:14:12 ISR1[2]: N: u2d7 SVD_DONE: Command Error = 0x3
Oct 26 17:14:12 ISR1[2]: N: u2d7 sid 0 stype 405 disk error 3
Oct 26 17:14:12 WXFT[2]: W: u2d7: Failed
Oct 26 17:14:12 WXFT[2]: W: u2d7 hard err in vol (v1) starting auto disable
```

```

:/: fru stat

```

CTLR	STATUS	STATE	ROLE	PARTNER	TEMP			
ulctr	ready	enabled	master	u2ctr	34.5			
u2ctr	ready	enabled	alt master	ulctr	33.5			

DISK	STATUS	STATE	ROLE	PORT1	PORT2	TEMP	VOLUME
uld1	ready	enabled	data disk	ready	ready	33	v0
uld2	ready	enabled	data disk	ready	ready	30	v0
uld3	ready	enabled	data disk	ready	ready	29	v0
uld4	ready	enabled	data disk	ready	ready	27	v0
uld5	ready	enabled	data disk	ready	ready	24	v0
uld6	ready	enabled	data disk	ready	ready	26	v0
uld7	ready	enabled	data disk	ready	ready	25	v0
uld8	ready	enabled	data disk	ready	ready	31	v0
uld9	ready	enabled	data disk	ready	ready	34	v0
u2d1	ready	enabled	data disk	ready	ready	30	v1
u2d2	ready	enabled	data disk	ready	ready	31	v1
u2d3	ready	enabled	data disk	ready	ready	30	v1
u2d4	ready	enabled	data disk	ready	ready	26	v1
u2d5	ready	enabled	data disk	ready	ready	34	v1
u2d6	ready	enabled	data disk	ready	ready	26	v1
u2d7	fault	disabled	data disk	bypass	bypass	-	v1
u2d8	ready	enabled	data disk	ready	ready	32	v1
u2d9	ready	enabled	data disk	ready	ready	26	v1

LOOP	STATUS	STATE	MODE	CABLE1	CABLE2	TEMP		
u2l1	ready	enabled	master	installed	-	31.0		
u2l2	ready	enabled	slave	installed	-	34.5		
u1l1	ready	enabled	master	-	installed	31.0		
u1l2	ready	enabled	slave	-	installed	35.0		

POWER	STATUS	STATE	SOURCE	OUTPUT	BATTERY	TEMP	FAN1	FAN2
ulpcu1	ready	enabled	line	normal	normal	normal	normal	normal
ulpcu2	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu2	ready	enabled	line	normal	normal	normal	normal	normal

Replacing FRUs

This section describes how to verify and confirm installation of a replacement FRU. It includes the following sections:

- “To Verify the FRU Firmware Level” on page 119
- “To Install the Replacement FRU” on page 121
- “To Verify the FRU Replacement” on page 124

▼ To Verify the FRU Firmware Level

The array contains various types of FRU firmware. The Sun StorEdge T3+ array contains the following firmware:

- Controller firmware:
 - Level 1
 - Level 2
 - Level 3
- Disk drive firmware
- Interconnect card firmware

When you install the new FRU, you will check its firmware level to determine if you need to upgrade the existing firmware. Use this procedure to check the firmware versions on your particular array model.

1. Type `ver` to display the header information:

- An example of header information for a Sun StorEdge T3+ array:

```
:/:<5>ver  
  
T3B Release 2.00 2001/08/02 15:21:29 (192.168.209.243)  
Copyright (C) 1997-2001 Sun Microsystems, Inc.  
All Rights Reserved.
```

In this example, the Sun StorEdge T3+ array is shown running version 2.00 level 3 controller firmware.

2. Type `fru list` to display the firmware versions for the disk drives and interconnect card on both array models.

For example:

```

:/: fru list

```

ID	TYPE	VENDOR	MODEL	REVISION	SERIAL
ulctr	controller card	0034	501-5710-02(0200	123456
u2ctr	controller card	0034	501-5710-02(0200	123455
uld1	disk drive	SEAGATE	ST118202FSUN	9329	LK478728
uld2	disk drive	SEAGATE	ST118202FSUN	9329	LK493799
uld3	disk drive	SEAGATE	ST118202FSUN	9329	LK493800
uld4	disk drive	SEAGATE	ST118202FSUN	9329	LK494457
uld5	disk drive	SEAGATE	ST118202FSUN	9329	NK040486
uld6	disk drive	SEAGATE	ST118202FSUN	9329	LK339351
uld7	disk drive	SEAGATE	ST118202FSUN	9329	LK150715
uld8	disk drive	SEAGATE	ST118202FSUN	9329	NK040157
uld9	disk drive	SEAGATE	ST118202FSUN	9329	LKG79907
u2d1	disk drive	SEAGATE	ST118202FSUN	9329	LKG90019
u2d2	disk drive	SEAGATE	ST118202FSUN	9329	LKH18597
u2d3	disk drive	SEAGATE	ST118202FSUN	9329	LKH15606
u2d4	disk drive	SEAGATE	ST118202FSUN	9329	LKH16563
u2d5	disk drive	SEAGATE	ST118202FSUN	9329	LKG88883
u2d6	disk drive	SEAGATE	ST118202FSUN	9329	LKH61907
u2d7	disk drive	SEAGATE	ST118202FSUN	9329	LKG90719
u2d8	disk drive	SEAGATE	ST118202FSUN	9329	LKG95442
u2d9	disk drive	SEAGATE	ST118202FSUN	EA29	LKG61774
u1l1	loop card	SCI-SJ	375-0085-01-	5.02 Flash	007924
u1l2	loop card	SCI-SJ	375-0085-01-	5.02 Flash	007382
u2l1	loop card	SCI-SJ	375-0085-01-	5.02 Flash	003343
u2l2	loop card	SCI-SJ	375-0085-01-	5.02 Flash	003344
ulpcu1	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001455
ulpcu2	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001408
u2pcu1	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001686
u2pcu2	power/cooling unit	TECTROL-CAN	300-1454-01(0000	001445
ulmpn	mid plane	SCI-SJ	375-0084-01-	0000	000650
u2mpn	mid plane	SCI-SJ	375-0084-01-	0000	000649

In this example:

- Disk drive firmware version is listed as Revision 9329
- Interconnect card (loop card) firmware version is listed as Revision 5.02 Flash
- Sun StorEdge T3 array controller EPROM firmware is listed as Revision 0200.
- Sun StorEdge T3+ array level 2 controller firmware, the level 2 image is 0200.

▼ To Install the Replacement FRU



Caution – Make sure you have the replacement FRU available before removing the failed FRU. If you remove a FRU, it must be immediately replaced. There is a timer in the unit that will cause the array and all connected arrays to power off in 30 minutes if a FRU is removed and not replaced.

- **Refer to the *Sun StorEdge T3+ Array Installation and Configuration Manual* for instructions on replacing FRUs.**

▼ To Rebuild a Replaced Drive FRU

A replaced drive FRU is automatically reconstructed by the array. If you should need to rebuild the replaced drive FRU manually, do so as follows:

1. **On the array, type:**

```
:/: vol recon undn
```

where *undn* is the unit (u) number (*n*) and the drive (d) number (*n*) of the drive you are rebuilding.

2. **Start a second telnet session with the array to check rebuild progress.**
3. **Check rebuild progress using the `proc list` command.**

The information in the `PERCENT` column and the `TIME` column, which shows the elapsed time, enables you to estimate when the volume will complete reconstruction.

```
:/: proc list

VOLUME          CMD_REF  PERCENT  TIME  COMMAND
v1              20241   23      0:09  vol recon
```

4. Check the drive status to ensure that the reconstruction of the replaced drive FRU has completed.

```

:/: fru stat

```

CTLR	STATUS	STATE	ROLE	PARTNER	TEMP			
ulctr	ready	enabled	master	u2ctr	32.0			
u2ctr	ready	enabled	alt master	ulctr	29.0			
DISK	STATUS	STATE	ROLE	PORT1	PORT2	TEMP	VOLUME	
uld1	ready	enabled	data disk	ready	ready	34	v0	
uld2	ready	enabled	data disk	ready	ready	34	v0	
uld3	ready	enabled	data disk	ready	ready	31	v0	
uld4	ready	enabled	data disk	ready	ready	34	v0	
uld5	ready	enabled	data disk	ready	ready	32	v0	
uld6	ready	enabled	data disk	ready	ready	38	v0	
uld7	ready	enabled	data disk	ready	ready	33	v0	
uld8	ready	enabled	data disk	ready	ready	34	v0	
uld9	ready	enabled	data disk	ready	ready	36	v0	
u2d1	ready	enabled	data disk	ready	ready	30	v1	
u2d2	ready	enabled	data disk	ready	ready	31	v1	
u2d3	ready	enabled	data disk	ready	ready	34	v1	
u2d4	ready	enabled	data disk	ready	ready	34	v1	
u2d5	ready	enabled	data disk	ready	ready	32	v1	
u2d6	ready	enabled	data disk	ready	ready	37	v1	
u2d7	ready	enabled	data disk	ready	ready	38	v1	
u2d8	ready	enabled	data disk	ready	ready	42	v1	
u2d9	ready	enabled	data disk	ready	ready	34	v1	
LOOP	STATUS	STATE	MODE	CABLE1	CABLE2	TEMP		
u2l1	ready	enabled	master	installed	-	30.0		
u2l2	ready	enabled	slave	installed	-	30.5		
u1l1	ready	enabled	master	-	installed	27.0		
u1l2	ready	enabled	slave	-	installed	29.5		
POWER	STATUS	STATE	SOURCE	OUTPUT	BATTERY	TEMP	FAN1	FAN2
ulpcu1	ready	enabled	line	normal	normal	normal	normal	normal
ulpcu2	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu2	ready	enabled	line	normal	normal	normal	normal	normal

Note – A delay can occur between the time when drive reconstruction is started and when the `fru stat` output displays an optimal status. If the drive is under reconstruction when you run the `fru stat` command, the command output will show the drive as missing. After the drive finishes rebuilding, the `fru stat` output shows the drive as ready.

Controller Card

Replacing a controller card does not require any special command line administration, other than verifying its state using the `fru stat` command. For information on the controller card and how to replace it, refer to the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Power and Cooling Unit

The replacement of the power cooling unit FRU should be done as quickly as possible to maintain full redundancy. For additional information including replacement instructions, refer to the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Note – The UPS battery, which is built into the power and cooling unit, will show a fault when the unit is first installed before it recharges.

Interconnect Card

Replacing an interconnect card does not require any special command line administration, other than verifying its state using the `fru stat` command. For information on the interconnect card and how to replace it, refer to the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Note – In `fru stat` output, information on the interconnect card is listed under the `LOOP` category.

Chassis and Centerplane

It is unlikely that the array chassis and centerplane, which is one FRU, will ever need to be replaced. If it does need to be replaced, contact your Sun service representative. Only trained, qualified service providers should replace an array chassis and centerplane.

▼ To Verify the FRU Replacement

1. **After replacing the FRU, use the `fru list` command to determine if the new FRU has the same firmware level as the FRU you replaced.**

See “To Verify the FRU Firmware Level” on page 119. If the firmware level of the new FRU is a later release, you will need to upgrade all firmware. For instructions on how to upgrade firmware, refer to the Service chapter in the *Sun StorEdge T3+ Array Installation and Configuration Manual*.

Note – If you replace a controller card FRU, the controller firmware automatically checks the firmware information on the system disk and upgrades or downgrades itself to the previous firmware version recorded. A manual controller firmware upgrade is not required.

2. Check the FRU status to verify that the failure has been corrected using the `fru stat` command.

```

:/: fru stat

```

CTLR	STATUS	STATE	ROLE	PARTNER	TEMP			
ulctr	ready	enabled	master	u2ctr	32.0			
u2ctr	ready	enabled	alt master	ulctr	29.0			
DISK	STATUS	STATE	ROLE	PORT1	PORT2	TEMP	VOLUME	
uld1	ready	enabled	data disk	ready	ready	34	v0	
uld2	ready	enabled	data disk	ready	ready	34	v0	
uld3	ready	enabled	data disk	ready	ready	31	v0	
uld4	ready	enabled	data disk	ready	ready	34	v0	
uld5	ready	enabled	data disk	ready	ready	32	v0	
uld6	ready	enabled	data disk	ready	ready	38	v0	
uld7	ready	enabled	data disk	ready	ready	33	v0	
uld8	ready	enabled	data disk	ready	ready	34	v0	
uld9	ready	enabled	data disk	ready	ready	36	v0	
u2d1	ready	enabled	data disk	ready	ready	30	v1	
u2d2	ready	enabled	data disk	ready	ready	31	v1	
u2d3	ready	enabled	data disk	ready	ready	34	v1	
u2d4	ready	enabled	data disk	ready	ready	34	v1	
u2d5	ready	enabled	data disk	ready	ready	32	v1	
u2d6	ready	enabled	data disk	ready	ready	37	v1	
u2d7	ready	enabled	data disk	ready	ready	38	v1	
u2d8	ready	enabled	data disk	ready	ready	42	v1	
u2d9	ready	enabled	data disk	ready	ready	34	v1	
LOOP	STATUS	STATE	MODE	CABLE1	CABLE2	TEMP		
u2l1	ready	enabled	master	installed	-	30.0		
u2l2	ready	enabled	slave	installed	-	30.5		
u1l1	ready	enabled	master	-	installed	27.0		
u1l2	ready	enabled	slave	-	installed	29.5		
POWER	STATUS	STATE	SOURCE	OUTPUT	BATTERY	TEMP	FAN1	FAN2
ulpcu1	ready	enabled	line	normal	normal	normal	normal	normal
ulpcu2	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu1	ready	enabled	line	normal	normal	normal	normal	normal
u2pcu2	ready	enabled	line	normal	normal	normal	normal	normal

Note – When checking disk drive status, a delay can occur between the time when drive reconstruction is started and the `fru stat` output displays an optimal status. If the drive is under reconstruction when you run the `fru stat` command, the command output shows the drive as missing. After the drive finishes rebuilding, the `fru stat` output shows the drive as ready.

▼ To Identify Miscabled Partner Groups

If you configure a partner group that has booted successfully; but you are unable to establish a telnet connection, a possible cause is that the partner group has been cabled together incorrectly.

The interconnect cable connections between dual controller units are critical for determining which unit is the master controller unit and which is the alternate master controller unit. If the interconnect cables are not properly installed on the interconnect cards, it is possible for the top unit to boot as the master controller, and the bottom unit would assume alternate master status. This becomes a problem because the host has been configured to use the MAC address of the bottom unit.

If the bottom unit is incorrectly cabled (making the bottom unit the alternate master), the bottom unit's Ethernet port will be inactive unless a failover situation occurs. In that event, the IP and MAC address of the bottom unit will take over the values of the master (top) unit.

If you suspect the partner group has been cabled together incorrectly, the following procedure can help you determine if the top unit is acting as the master controller.

1. Determine the MAC address of the top unit.

The MAC address is located on a pull-out tab at the front of the unit, to the left of the first disk drive as shown in FIGURE 11-1.

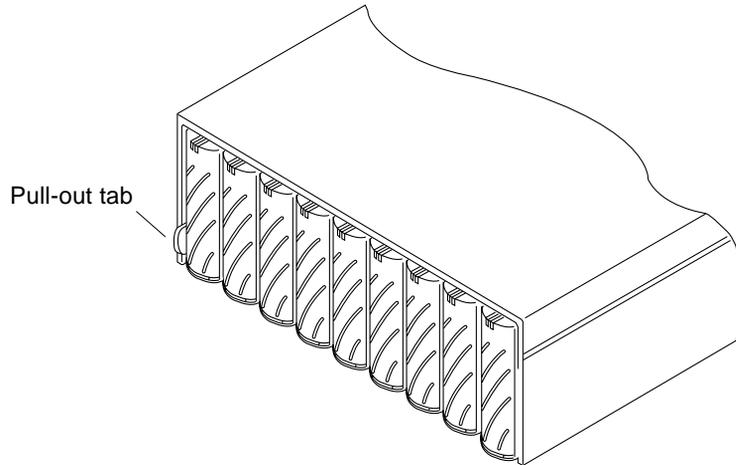


FIGURE 11-1 MAC Address on Pull-out Tab

2. Edit the files on the RARP server to include the MAC address of the top unit.

a. Edit the `/etc/ethers` file by adding the MAC address and array name.

For example:

```
8:0:20:7d:93:7e array-name
```

In this example:

- `8:0:20:7d:93:7e` is the MAC address
- `array-name` is the name of the master controller unit.

b. Edit the `/etc/hosts` file with the IP address and array name.

For example:

```
192.129.122.111 array-name
```

In this example, `192.129.122.111` is the assigned IP address.

c. Edit the `/etc/nsswitch.conf` file to reference the local system files.

To ensure the Solaris software environment uses the changes made to `/etc/ethers` and `/etc/hosts` files, edit the `host` and `ethers` entries in the `/etc/nsswitch.conf` file so that the `files` parameter appears before the `[NOTFOUND=return]` statements as shown:

```
hosts:      nis files [NOTFOUND=return]
ethers:     nis files [NOTFOUND=return]
```

d. Determine if the `rarp` daemon is running by typing:

```
# ps -eaf | grep rarpd
```

- If the `rarp` daemon is running, you must kill the process and restart it so the host to recognize the changes you just made. To kill the process, type:

```
# kill process ID
```

The process ID was provided in the `ps -eaf | grep rarpd` command output.

- If the `rarp` daemon is not running, continue to the next step.

e. Start the `rarp` daemon in the Solaris environment by typing:

```
# /usr/sbin/in.rarpd -a &
```

3. Make sure you have an Ethernet connection to the Ethernet port of the top unit.

See the cabling section in the *Sun StorEdge T3+ Array Installation and Configuration Manual* for more information. If you do not have an Ethernet connection to the top unit but do have one to the bottom unit, remove the cable from the bottom unit and connect it to the top unit.

4. Press the power switch on the power and cooling units on both arrays to remove AC power (FIGURE 11-2).

After pressing the switches, it may take some time for the units to power off while shutdown procedures are performed. Wait until the units have powered off completely.

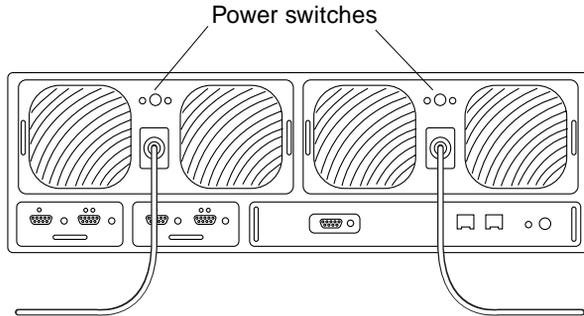


FIGURE 11-2 Power Switch Locations

5. **After both units have powered off, press the power switch on the power and cooling units again to restore power to and reset the arrays.**

It may take up to several minutes for the arrays to power on and come back online. All LEDs will be green when the unit is fully powered on.

6. **After the units are fully powered on, start a telnet session.**

The telnet session will connect to the top unit.

If you still cannot start a telnet session, investigate the following other possible causes:

- **rarp server is not responding.** To determine if this is the problem:
 - Verify that the RARP daemon is running on the host system.
 - Verify that the `/etc/nsswitch.conf` file is properly configured on the RARP server.
 - In the Solaris environment, use the `snoop(1M)` command to verify that the array is attempting to establish RARP communication with the Solaris server.
- **MAC address is incorrect.** In the Solaris operating environment, use the `snoop(1M)` command to specify the MAC address of the array and see if any RARP packets are transmitted. If nothing is observed during a reboot of the array, verify that the MAC address on the array label matches the MAC address configured on the RARP server.
- **Netmask is incorrect.** The default netmask address used on the array is 255.255.255.0. If your local subnet uses a different netmask, the RARP operation may not work.
- **Inoperable network connections.** If you are using hubs to connect to the network, try eliminating or replacing the hub.

Identifying Data Channel Failures

The data channel encompasses the host data path that extends from the host bus adapter to the array. Errors in the host data channel are outside of the scope of the Sun StorEdge T3+ array. To determine failures in the data path, you must use host-based application diagnostics, such as the StorTools product for the Solaris operating environment.

Refer to the documentation of the selected diagnostics tool for information on identifying data channel failures.

CLI Commands

This chapter contains descriptions of the commands supported by the Sun StorEdge T3+ array.

Topics in this chapter include:

- “Complete List of CLI Commands” on page 132
- “FRU Identifiers” on page 135
- “Code Conventions” on page 135
- “General Commands” on page 136
- “Volume slicing, Mapping, and Masking Commands” on page 158
- “Fibre Channel Topology Command” on page 185
- “Network Time Protocol Commands” on page 189

Complete List of CLI Commands

TABLE 12-1 lists all the commands described in this chapter. This includes all the general T3+ commands, and those commands common to controller upgrades with firmware version 2.1 when volume slicing is enabled. In the table, volume slicing, mapping, and masking commands are referred to as “Multi-LUN.”

TABLE 12-1 CLI Commands Listed in Alphabetical Order for T3+ Array with Controller Firmware Version 2.1

Command	Description	Command Type	Page
boot	Boot the system.	General	12-137
disable	Disable specified FRUs.	General	12-138
disk	Perform disk administration.	General	12-138
enable	Enable specified FRUs.	General	12-139
ep	Program the flash EPROM.	General	12-139
fru	Display the FRU information.	General	12-140
help	Display reference manual pages.	General	12-141
hwwn add	Add a WWN number to an existing group.	Multi-LUN	12-182
hwwn list	List all the WWN(s) in a specified group.	Multi-LUN	12-180
hwwn listgrp	List all the registered group names.	Multi-LUN	12-181
hwwn rm	Remove one or more WWN number(s) from a specific group.	Multi-LUN	12-183
hwwn rmgrp	Remove a specific group name and all its associated WWNs.	Multi-LUN	12-184
id	Display FRU identification summary.	General	12-142
logger	Generate messages to the syslog in the unit.	General	12-143
lpc	Get the interconnect card property.	General	12-144
lun default	Set up the LUN default permission.	Multi-LUN	12-176
lun map add	Maps a LUN to a slice number.	Multi-LUN	12-168

TABLE 12-1 CLI Commands Listed in Alphabetical Order for T3+ Array with Controller Firmware Version 2.1 *(Continued)*

Command	Description	Command Type	Page
lun map list	List all current LUN and slice mapping entries.	Multi-LUN	12-166
lun map rm	Remove the mapping between a specified LUN and its corresponding slice.	Multi-LUN	12-169
lun map rm all	Remove all current LUN mappings.	Multi-LUN	12-170
lun perm	Grant LUN access to: <ul style="list-style-type: none"> • A single HBA (specified by its WWN) • A group of HBAs • All HBAs 	Multi-LUN	12-174
lun perm list	List the accessibility of LUNs based on specified criteria.	Multi-LUN	12-172
lun wwn list	List all registered WWNs.	Multi-LUN	12-177
lun wwn rm	Remove the specified WWN from the system database.	Multi-LUN	12-178
lun wwn rm all	Remove all registered WWNs from the system database.	Multi-LUN	12-179
ntp	Display: <ul style="list-style-type: none"> • The NTP server IP address • The polling on/off status • The polling interval in minutes. 	Network Time Protocol	12-196
ntp interval	Set the polling interval from 1 to 60 minutes.	Network Time Protocol	12-199
ntp poll	Turn NTP polling on and off.	Network Time Protocol	12-198
ntp server	Establish the server that is running the NTP daemon in the T3 array firmware.	Network Time Protocol	12-200
ntp stats	Display: <ul style="list-style-type: none"> • The time of the last poll • The NTP server IP address • The adjustment made to the T3+ array internal clock • The clock set success statement 	Network Time Protocol	12-197

TABLE 12-1 CLI Commands Listed in Alphabetical Order for T3+ Array with Controller Firmware Version 2.1 *(Continued)*

Command	Description	Command Type	Page
ofdg	Perform back-end offline loop diagnostics.	General	12-145
port	Configure the interface port.	General	12-146
proc	Display status of outstanding vol processes.	General	12-147
refresh	Start/stop battery refreshing or display its status.	General	12-148
reset	Reset the system.	General	12-149
set	Display or modify the set information.	General	12-150
set timezone	Set the timezone of the server that is running the NTP daemon.	Network Time Protocol	12-190
shutdown	Shutdown an array or partner group.	General	12-152
sys	Display or modify the system information.	General	12-153
sys enable_volslice	Enable volume slicing.	Multi-LUN	12-159
sys fc_topology	Set the T3+ array topology firmware mode.	Fibre Channel Topology	12-186
ver	Display the software version.	General	12-154
vol	Display or modify the volume information.	General	12-155
volslice create	Create a volume slice.	Multi-LUN	12-162
volslice list	List the details of all or specified volume(s).	Multi-LUN	12-160
volslice remove	Remove a volume slice or slices.	Multi-LUN	12-164

FRU Identifiers

Many commands use a FRU identifier to refer to a particular FRU in an array. This identifier contains a unit constant (*u*), the unit number (*encid*), the FRU constant (*ctr* for controller card, *pcu* for power and cooling unit, *1* for interconnect card, *d* for disk drive) and the FRU number (*n*). TABLE 12-2 lists the possible FRU variables as they appear in this chapter.

TABLE 12-2 FRU Identifiers

FRU	Identifier	Unit Number
Controller card	<code>uencidctr</code>	<i>encid</i> = unit number (1, 2)
Power and cooling unit	<code>uencidpcu[1 2]</code>	<i>encid</i> = unit number (1, 2) <i>n</i> = pcu number (1, 2)
Interconnect card	<code>uencid1[1 2]</code>	<i>encid</i> = unit number (1, 2) <i>n</i> = interconnect card number (1, 2)
Disk drive	<code>uencidn</code>	<i>encid</i> = unit number (1, 2) <i>n</i> = disk drive number (1 to 9)

Code Conventions

This chapter contains the following typographic conventions:

TABLE 12-3 Text Code Conventions

<i>Italic font</i>	Indicates options or parameters that you replace with an appropriate word or string.
Boldface courier font	Indicates optional or mandatory information that is typed verbatim.
[] Brackets	Indicates optional arguments.
Pipe	Separates mutually exclusive arguments.

General Commands

This section describes the commands supported by every Sun StorEdge T3+ array. For additional command information, use the *command-name help* command.

For example, for information on the `reset` command, type:

```
:/: reset help
usage: reset [ -y ]
```

This section describes the following commands:

- “boot” on page 137
- “disable” on page 138
- “disk” on page 138
- “enable” on page 139
- “ep” on page 139
- “fru” on page 140
- “help” on page 141
- “id” on page 142
- “logger” on page 143
- “lpc” on page 144
- “ofdg” on page 145
- “port” on page 146
- “proc” on page 147
- “refresh” on page 148
- “reset” on page 149
- “set” on page 150
- “shutdown” on page 152
- “sys” on page 153
- “ver” on page 154
- “vol” on page 155

General Command Descriptions

This section provides summaries of Sun StorEdge T3+ array commands. See “FRU Identifiers” on page 135 and “Code Conventions” on page 135 for help in understanding the information. Command summaries are listed in alphabetical order.

boot

Synopsis

`boot -f disk-no.`

`boot -i filename.bin`

`boot -s`

`boot -w`

Function

Boots the RAID software according to the boot mode configured with the `set` command.

Parameters

TABLE 12-4 `boot` Command Options and Parameters

Parameter	Function
<code>-f disk-no.</code>	Verifies the disk boot image in the specified disk; updates the disk with its current disk position if the boot partition is OK; if the boot partition is not OK, the command exits, displaying “Failed to set new boot disk position” on the console.
<code>-i filename.bin</code>	Copies the level 3 controller firmware image to the controller card flash. The <i>filename</i> string can be a maximum of 11 characters.
<code>-s</code>	Used when disks in an array are detected as being in the wrong disk position. If this occurs, a special prompt ADMIN> is displayed. At this point, use the <code>boot -s</code> command to correct this state and reboot the array.
<code>-w</code>	Sets the default values to system area and configuration setting only. The boot image and local file system will be retained. Use to quickly clear all the system and configuration settings.

disable

Synopsis

`disable uencid`

Function

Disables specified FRUs.

Parameters

TABLE 12-5 `disable` Command Options and Parameters

Parameter	Function
<i>uencid</i>	See “FRU Identifiers” on page 135.

disk

Synopsis

`disk version drives`

`disk download drives filename`

Function

Controls, displays, and configures disks.

Parameters

TABLE 12-6 `disk` Command Options and Parameters

Parameter	Function
<code>version</code>	Displays the firmware version of the selected drives(s).
<code>download</code>	Updates the selected drives(s) with new firmware specified by the <i>filename</i> .
<i>drives</i>	Specifies the disk drive numbers.
<i>filename</i>	Specifies the file to use for drive firmware updates.

enable

Synopsis

`enable uencid`

`enable uencid1 [1 | 2]`

Function

Enables specific FRUs.

Parameters

TABLE 12-7 `enable` Command Options and Parameters

Parameter	Function
<code>uencid, 1</code>	See “FRU Identifiers” on page 135.

ep

Synopsis

`ep download filename`

Function

Updates the FLASH with new firmware specified by *filename*.

Parameters

TABLE 12-8 `ep` Command Options and Parameters

Parameter	Function
<i>filename</i>	Specifies the new firmware to be loaded into the controller EPROM. The file must reside in the local file system.

fru

Synopsis

```
fru l[ist]
fru list sys
fru list [u<1-8>]
fru list [u<1-8>]c[1]
fru list [u<1-8>]d[<1-9>]
fru list [u<1-8>]l[<1-2>]
fru list [u<1-8>]pcu[<1-2>]
fru list [u<1-8>]mpn[1]
fru s[tat]
fru stat sys
fru stat [u<1-8>]
fru stat [u<1-8>]c[1]
fru stat [u<1-8>]d[<1-9>]
fru stat [u<1-8>]l[<1-2>]
fru stat [u<1-8>]pcu[<1-2>]
fru st[atistic]
fru statistic sys
fru statistic [u<1-8>]
fru statistic [u<1-8>]d[<1-9>]
fru statistic [u<1-8>]p[1] [clear]
fru statistic [u<1-8>]v[<1-2>] [clear]
fru myuid
```

Function

Displays the field-replaceable unit (FRU) information.

Parameters

TABLE 12-9 fru Command Options and Parameters

Parameter	Function
<code>list</code> <code>list sys</code>	Displays all FRU ID type, vendor, model, version, and serial number.
<code>stat sys</code>	Displays the status and role of each T3+ array.
<code>statistic sys</code>	Displays I/O statistics.
<code>myuid</code>	Displays the current controller.
<code>uencid, d, l, pcu</code>	See “FRU Identifiers” on page 135.
<code>c[1]</code>	Controller card.
<code>mpn[1]</code>	Midplane card.
<code>p[1]</code>	Controller card port.
<code>v[<1-2>]</code>	Volume.
<code>clear</code>	Resets the counter to restart calculating the statistics.

help

Synopsis

command help

Function

Displays commands with help text explanations.

The `help` command is a *re-entrant command*. A re-entrant command is one that can be run in multiple telnet sessions (CLI windows) to the same array.

Parameters

TABLE 12-10 help Command Options and Parameters

Parameter	Function
<i>command</i>	Any valid T3+ array CLI command.

id

Synopsis

`id read uencid`

`id read uencidpcu[1 | 2]`

`id read uencid1[1 | 2]`

`id read uencidc1`

Function

Displays the FRU identification summary.

Parameters

TABLE 12-11 `id` Command Options and Parameters

Parameter	Function
<code>uencid</code> , <code>pcu</code> , <code>1</code> , <code>c</code>	See “FRU Identifiers” on page 135..
<code>c1</code>	Controller card.

logger

Synopsis

```
logger [-dmprstlog] [-f <file | ->] [-p facility.priority] [message]
```

Function

Generates messages to the syslog in the unit.

Parameters

TABLE 12-12 logger Command Options and Parameters

Parameter	Function
-dmprstlog	Dumps the reset log, and displays the system crash information. This log information can be useful in troubleshooting controller crashes.
-f	Specifies a particular file or content to be logged. -f <file> specifies a particular file; if used, only the first 20 lines are logged. -f <-> specifies a request to log standard input; if used, the next line up to Control-c or Control-d followed by a Return is logged.
-p	Specifies the message's facility and priority level. The default facility is user and the default priority is notice.
message	Text of message to be logged.

Example

```
:/: logger -p syslog.warning this is a test
```

This example shows text being written to the `syslog` facility as a warning and writes the following entry into the `syslog`:

```
Oct 26 14:44:56 sh01[1]: W: this is a test
```

lpc

Synopsis

```
lpc version
lpc download uencid1[1 | 2] filename
lpc reboot uencid1[1 | 2]
lpc rsn uencid1[1 | 2]
lpc rid uencid1[1 | 2]
lpc leadtest uencid1[1 | 2] delay
```

Function

Used to display and upgrade interconnect cards.

Parameters

TABLE 12-13 `lpc` Command Options and Parameters

Parameter	Function
<code>version</code>	Displays each interconnect card's firmware version.
<code>download</code>	Downloads new interconnect card firmware.
<code>reboot</code>	Executes soft reset to the selected interconnect card.
<code>rsn</code>	Displays the serial number.
<code>rid</code>	Displays interconnect card's vendor id and product id.
<code>ledtest</code>	Tests LEDs controlled through the interconnect card. This option does not test the AC LED on the PCU or the drive activity LED on disk drives; these LEDs are not controlled through the interconnect card.
<code>uencid, 1</code>	See "FRU Identifiers" on page 135.
<code>filename</code>	File name of new firmware to download.
<code>delay</code>	Specifies the delay in seconds between LED pattern changes.

ofdg

Synopsis

```
ofdg [-y] health_check  
ofdg [-y] fast_test uencid1[1 | 2]  
ofdg [-y] fast_find uencid1[1 | 2]  
ofdg [-y] find uencid1[1 | 2]
```

Function

A utility to perform back-end loop diagnostics. All non-u1 disks must be assigned to a volume to perform this diagnostic. Test output is displayed in the syslog file. This test does not detect missing interconnect cables.

Parameters

TABLE 12-14 ofdg Command Options and Parameters

Parameter	Function
<code>health_check</code>	Performs a fast test on all detected loops in a partner group.
<code>fast_test</code>	Performs a nondisruptive, fast loop test on the specified loop. Disk path failovers are not performed with this test.
<code>fast_find</code>	Performs a loop test on partner groups, providing limited detection of bad FRUs. Use the <code>fast_find</code> option before using the <code>find</code> option to eliminate interconnect cards, interconnect cables, and controller cards as possible bad FRUs. This test reconfigures loops and performs disk path failovers.
<code>find</code>	Performs a loop test on all loops in a partner group and provides loop fault diagnostics, which have the capability to isolate a FRU problem to a single disk port. Use the <code>find</code> option only after using the <code>fast_find</code> option. This test reconfigures loops and by-passes disk ports.
<code>-y</code>	Provides an affirmative answer to the command prompt: WARNING - Volume data will be offline while OFDG is running. Continue? [N]:
<code>uencid, 1</code>	See “FRU Identifiers” on page 135.

port

Synopsis

```
port list
```

```
port set port targetid target-id-value | s[oft]
```

```
port host port sun | other
```

```
port listmap [port]
```

Function

This is a port management utility.

Parameters

TABLE 12-15 port Command Options and Parameters

Parameter	Function
list	Displays the configuration of the ports.
set	Modifies port parameters and changes the port target ID value. Any changes to the target ID require a reset for the changes to take effect.
host	Whenever attaching to a Solaris 2.6 host, add the option 'sun' to the end of the <code>port host</code> command. For example: <code>port host ulp1</code>
listmap	Shows the current mappings between host interface ports and LUNS and volumes.
port	Specifies the port. The format is: <code>uarraypportindex</code> . Examples: <code>ulp1</code> , <code>u2p1</code>
target-id-value	Specifies the target ID (FC-AL ALPA). Valid values are from 0 to 125; the optional argument <code>s</code> means soft-addressing.

Example

```
:/: port list
port targetid addr_type status host wwn
ulp1 1 hard online other 20020f230000036a
```

```

:: port listmap
port  targetid  addr_type  lun  volume  owner  access
ulp1  1          hard       1   ulvoll  ul     primary

```

```

:: port set ulp1 targetid 6

```

proc

Synopsis

```
proc list [cmd-ref | vol-name]
```

```
proc kill cmd-ref | vol-name
```

Function

Displays the status of outstanding vol processes.

Parameters

TABLE 12-16 proc Command Options and Parameters

Parameter	Function
list	Displays the command reference number of an active vol <code>init</code> , <code>vol verify</code> or <code>vol recon</code> command currently in progress.
kill	Terminates the active vol <code>init</code> , <code>vol verify</code> or <code>vol recon</code> command by specifying its command reference number.
<i>cmd-ref</i>	Specifies the command reference number.
<i>vol-name</i>	Specifies the volume name.

refresh

Synopsis

refresh -c

refresh -s

refresh -i

refresh -k

Function

Starts the battery refresh operation on all power and cooling unit 1 batteries, followed by all power and cooling unit 2 batteries. This operation consists of discharging and recharging each power and cooling unit battery. Battery refresh takes up to 10 hours for each power and cooling unit.

An automatic battery refresh cycle is performed once every 28 days; the cycle frequency is specified in the `/etc/sched.conf` file on the array. A manual battery refresh cycle is not required. During this cycle, one battery is always fully charged. A battery refresh cycle will not begin unless both batteries are in optimal condition.

Parameters

TABLE 12-17 `refresh` Command Options and Parameters

Parameter	Function
<code>-c</code>	Starts the refresh cycle.
<code>-s</code>	Reports the refresh cycle status. Note that in the output, the date for the “Last Refresh” is actually that of the last refresh <i>attempt</i> . Therefore, the refresh success, or lack of success, is not reported at this time.
<code>-i</code>	Directs the scheduler to reload the battery configuration file; for field service use only.
<code>-k</code>	Kills the current refreshing task.

reset

Synopsis

reset [-y]

Function

Resets the controller. Use the `sync` command before a `reset` command.

The `reset` command is a *re-entrant command*. A re-entrant command is one that can be run in multiple telnet sessions (CLI windows) to the same array.

Parameters

TABLE 12-18 `reset` Command Options and Parameters

Parameter	Function
-y	Answers “yes” when prompted to perform reset.

set

Synopsis

```
set bootmode auto | tftp | none
```

```
set bootdelay seconds
```

```
set ip value
```

```
set netmask value
```

```
set gateway value
```

```
set tftphost value
```

```
set tftpfile value
```

```
set hostname value
```

```
set logto 1 | * | /filename
```

```
set loglevel 0 | 1 | 2 | 3 | 4
```

```
set rarp on | off
```

Function

Sets and displays the boot parameters. Without any parameters, the set command shows the values of the parameters.

Parameters

TABLE 12-19 set Command Options and Parameters

Parameter	Function
bootmode	Defines how the array is booted: <ul style="list-style-type: none">• Use <code>auto</code> boot for loading the RAID software from the local drives. This is the default mode.• If <code>bootmode</code> is using <code>tftp</code>, use <code>tftphost</code> and <code>tftpfile</code>. (Field service procedure only.)• If <code>bootmode</code> is set to <code>none</code>, use system EPROM command line. (Field service procedure only.).
bootdelay <i>seconds</i>	Specifies the delay time, in seconds, to cancel a system boot. Use value from 1 to 20 seconds.
ip	Specifies the network IP address. The <i>value</i> is expressed in dotted decimal notation (for example, 192.9.200.100).

TABLE 12-19 `set` Command Options and Parameters (Continued)

Parameter	Function
<code>netmask</code>	The subnet mask for the network. The <i>value</i> is expressed in dotted decimal notation (for example, 255.255.255.0).
<code>gateway</code>	The default gateway on the network. The <i>value</i> is expressed in dotted decimal notation (for example, 192.9.200.1).
<code>tftpboot</code>	The IP address of the server from which a <code>tftp boot</code> should be performed. (Field service procedure only.)
<code>tftpfile</code>	The file name on the tftpboot server that contains the RAID software. The <i>value</i> is a string of up to 39 characters. (Field service procedure only.)
<code>hostname</code>	The user-assigned name of the array.
<code>logto</code>	<p>The file where messages are to be logged:</p> <ul style="list-style-type: none"> • If <code>logto</code> is set to 1, logging occurs only through the serial port; log messages are not sent to the <code>/syslog</code> file. • If <code>logto</code> is set to *, and an <code>/etc/syslog.conf</code> entry specifies an output file name that already exists, the array appends new log information to the end of that output file. If the output file specified in <code>/etc/syslog.conf</code> doesn't exist, the array creates a new file with that name. If the output file does not exist and cannot be created, the entry will be ignored. <p>Note: If there are changes to the <code>/etc/syslog.conf</code> file, you must type <code>set logto *</code> for those changes to take effect.</p> <ul style="list-style-type: none"> • If <code>logto</code> is set to <code>/filename</code>, and the specified file name does not exist, the array creates a new file with that name. If the file name already exists, the array appends new log information to the end of the existing file. If the file name does not exist and cannot be created, the array will attempt to open or create the default syslog file <code>/syslog</code>. If both <code>/filename</code> and <code>/syslog</code> cannot be opened, logging occurs through the serial port only.
<code>loglevel</code>	The level of messages to be logged (0=none, 1=errors, 2=warnings/errors, 3=notices/warnings/errors, 4=informational/notices/warnings/errors).
<code>rarp</code>	If enabled and connected to a rarp server, sets the IP address.

shutdown

Synopsis

shutdown [-y]

Function

The shutdown command performs a controlled shutdown of the array and any array connected to it.

Parameters

TABLE 12-20 shutdown Command Options and Parameters

Parameter	Function
-y	Answers “yes” when prompted to perform a shutdown.

sys

Synopsis

`sys list`

`sys stat`

`sys blocksize 16k | 32k | 64k`

`sys cache auto | writebehind | writethrough | off`

`sys mirror auto | off`

`sys mp_support none | rw | mpxio | std`

`sys rd_ahead on | off`

`sys recon_rate high | med | low`

`sys naca`

Function

This utility manages system information.

Parameters

TABLE 12-21 `sys` Command Options and Parameters

Parameter	Function
<code>list</code>	Displays the current controller configuration.
<code>stat</code>	Displays the current controller status.
<code>blocksize</code>	Sets the system block size, which equals the stripe unit size: the amount of data written to each drive when striping data across drives. The block size can be changed only when there are no volumes defined.
<code>cache</code>	Sets the system caching policy.
<code>mirror</code>	Turns the cache mirror feature on or off when using dual controllers.

TABLE 12-21 `sys` Command Options and Parameters (Continued)

Parameter	Function
<code>mp_support</code>	<ul style="list-style-type: none">• If <code>mp_support</code> is set to <code>none</code>, multipathing support is disabled. Use this mode for workgroup configurations only.• If <code>mp_support</code> is set to <code>rw</code>, multipathing support is enabled. Use this mode for enterprise configurations to ensure redundancy.• If <code>mp_support</code> is set to <code>mpxio</code>, this setting supports Sun StorEdge Traffic Manager software's multipathing utility. This setting is required for all enterprise configurations in Sun Cluster environments.• The <code>std</code> option is reserved. Before using this option, contact your SunService™ provider.
<code>rd_ahead</code>	Sets the number of sequential read commands before automatically switching to datablock read ahead. Set to <code>off</code> to always perform datablock read ahead.
<code>recon_rate</code>	The disk drive reconstruction rate options: <code>high</code> , <code>medium</code> , or <code>low</code> .
<code>naca</code>	Enables IBM/AIX support.

`ver`

Synopsis

`ver`

Function

Displays the current version of level 3 controller firmware on a Sun StorEdge T3+ array.

vol

Synopsis

```
vol add name data drive raid 0 | 1 | 5 [standby drive]
vol disable drive [to-standby]
vol init name data | sysarea | fast [rate <1-16>]
vol list name
vol mode name
vol mount name
vol recon drive to-standby | from-standby
vol remove name
vol stat name
vol unmount name
vol verify name [fix] [rate <1-8>]
vol -y remove name
```

Function

This command manages volumes.

Parameters

TABLE 12-22 vol Command Options and Parameters

Parameter	Function
add	Configures a new volume by assigning it data and hot spare (standby) drives. The data drives cannot already be assigned to any other volume. The hot spare (standby) drive can be shared by multiple volumes.
disable	Marks the specified drive unusable. If the <code>to-standby</code> argument is included, data is first reconstructed onto the hot spare (standby) drive (if available), and then it is disabled. The volume must be mounted for this command.

TABLE 12-22 vol Command Options and Parameters (Continued)

Parameter	Function
init	Initializes the volume parity or system area label. This argument specifies the initialize option to be performed on the selected volume: <ul style="list-style-type: none">• data - Initializes the entire data area with correct parity data. This operation will destroy all user data on the volume. It will take several minutes to complete.• sysarea - Initializes the system area with correct private label information. This option is only for qualified service providers.• fast - Initializes and bypasses data area. The initialization rate can be set to any value from 1 (slowest) to 16 (fastest). The default rate for the array is 16. This option is to be used for RAID 0 volumes only.
list	Displays the configuration of the currently assigned volumes.
mode	Displays cache and mirror settings, and whether the volume is mounted.
mount	Mounts the volume so the host can access it.
recon	Rebuilds (reconstruct) the data of the specified drive. If the to-standby argument is included, data is written to the hot spare (standby) drive. If the from_standby argument is included, data is written back from the hot spare (standby). The volume must be mounted for this command. If necessary, the drive is spun up before reconstruction occurs. Following successful reconstruction, the disabled drive is re-enabled.
remove	Deletes a volume by unassigning the data and hot spare (standby) drives assigned to it. The drives can then be assigned to a different volume. The volume must be unmounted before it can be removed.
stat	Displays the operational status of a volume.
unmount	Unmounts the volume to disable access from the host.
verify	Enables an array administrator to execute manual parity checks on existing volumes. The default verification rate is 1 (minimum impact on data host). The verification rate can be set to a value from 1 to 8. Run the vol verify command before performing tape-backup overwrite cycles, approximately once every 30 days. If the fix option is specified and an error is detected, the vol verify command will regenerate parity from the existing data on the volume.
-y remove	Deletes a volume without prompting for verification to remove.
name	Specifies the volume name.
drives	Specifies the set of drives to be assigned to the volume.

Example

```
:/: vol list
volume    capacity    raid    data    standby
v0        125.2 GB    5       u1d1-8  u2d9
v1        125.2 GB    5       u2d1-8  u2d9
```

```
:/: vol add vol-name data u5d1-8 raid 5 standby u5d9
```

Volume slicing, Mapping, and Masking Commands

With firmware version 2.1, volume slicing, mapping, and masking operations are enabled with the `sys volslice_enable` command. For a discussion of the volume slicing, mapping, and masking features and limitations, see “Volume Slicing, Mapping, and Masking Configuration” on page 71.

This section covers the following topics:

- “Volume Slicing Command Descriptions” on page 158
- “LUN Mapping Command Descriptions” on page 165
- “LUN Masking Commands Descriptions” on page 171

Volume Slicing Command Descriptions

This section describes the following commands:

- “`sys enable_volslice`” on page 159
- “`volslice list`” on page 160
- “`volslice create`” on page 162
- “`volslice remove`” on page 164

sys enable_volslice

Synopsis

```
sys enable_volslice
```

Function

This command enables volume slicing. For this command to be successful with an enterprise configuration, *both* arrays must be running firmware version 2.1 and higher. To verify the controller firmware version, see *Sun StorEdge™ T3+ Array Installation and Configuration Manual*.

Sun StorEdge T3+ arrays might have one or two volumes set up. When you enable volume slicing, a default slice for each existing volume is created having the size of that volume. The resulting default slices are mapped to LUNs as 1:1 (one slice mapped to one LUN). The resulting default permissions are read/write.

Slicing a slice is not allowed. Therefore, if you want more than two slices, you must first remove the two default slices. To set up the Sun StorEdge T3+ array to do this and allow for volume slicing, mapping, and masking operations, see “To Enable Volume Slicing, Mapping, and Masking Operations” on page 76.



Caution – A procedure exists to disable the volume slicing feature; be aware that::

- The process is involved; detailed administrator knowledge is required.
- The risk of data loss exists.

If you still want to revert to the prior state before enabling volume slicing, see “To Enable Volume Slicing, Mapping, and Masking Operations” on page 76.

Example

```
# sys enable_volslice
WARNING - Once enabled this feature cannot be disabled
continue? Y
Initializing volslice configuration...
Volume slicing is enabled now!
2 out of Max. 16 slices created, 14 available. (example)
:::
```

volslice list

Synopsis

```
volslice list
```

```
volslice list slice_name
```

```
volslice list -v volume_name
```

Function

Lists the details of all or specified volume(s). If a slice name is provided, the details of that slice are displayed. If a volume name is provided, the details of the slices and unused portion on that volume are shown. With no options specified, the details of all volumes (slices and unused portion) are shown.

Parameters

TABLE 12-23 volslice list Command Options and Parameters

Parameter	Function
<i>slice_name</i>	Specifies the slice name such as s0, s1, s2, and s3.
-v <i>volume_name</i>	Specifies the volume name such as v0, v1, v2, and v3.

Example

- The first example lists the details of all slices (slices and unused portion) on all volumes.

```
# volslice list
Slice      Slice Num    Start Blk    Size Blks    Capacity    Volume
s0         0            0            104870400    50.005 GB   v1
-         -            104870400    31488        0.014 GB    v1
s1         1            0            104901888    50.020 GB   v2
-         -            104901888    140125327    62.3 GB     v2
```

In the above example, slice s0 has available space, starting with block 104,870,400 and extending for 31,488 blocks. However, since these 31,488 blocks equates to 0.014 Gbyte, it does not meet the 1-Gbyte slice size requirement and thus cannot be further sliced. The unused space (140,125,327 blocks) of volume v2, however, does allow for further slicing.

- The second example lists the details of slice s1.

```
# volslice list s1
Slice      Slice Num   Start blk   Size blks   volume
s1         1           0           104901888   v2
```

- The third example lists the details of the slices and unused portion on volume v1.

```
# volslice list -v v1
Slice      Slice Num   Start blk   Size blks   volume
s0         0           0           104870400   v1
-         -           104870400   31488       v1
-         -           104901888   87418240    v1
```

volslice create

Synopsis

```
volslice create
```

```
volslice create slice_name -z size volume_name
```

```
volslice create slice_name -s start -z size volume_name
```

Function

Creates a volume slice. A maximum of 16 slices is allowed for an array workgroup or an enterprise configuration.

Note – If you issue the `volslice create` command without any parameters, you will be interactively prompted to specify the parameters.

Note – The default permission of a newly created slice is *none*.

Parameters

TABLE 12-24 volslice create Command Options and Parameters

Parameter	Function
<i>slice_name</i>	Specifies the slice name such as s0, s1, s2, and s3. Can be any 12 (maximum) alphanumeric characters except a blank space.
-z <i>size</i>	Specifies the size of the block in Gbytes.
<i>volume_name</i>	Specifies the volume name such as v0, v1, v2, and v3.
-s <i>start</i>	In addition to the above options, this parameter can be used to specify the starting block. If it is not specified, the application will find the best fit for the specified slice size on the specified volume and create the slice accordingly. Use the <code>volslice list</code> command to find the starting block of unused space on a slice.

Example

This example shows the interactive and manual methods to create a slice.

```
# volslice create
volume      capacity      raid    data      standby
v1          50.020 GB      5      uld1-4    uld9
v2          50.020 GB      5      uld5-8    uld9
Select the volume to be sliced: v1
Slice       Slice Num      Start Blk    Size Blks    Capacity     Volume
-          -              0            143709696    50.020 GB   v1
Enter the name of slice to be created : s0
Do you want to specify the start block for slice ? [N]: n
Enter size of slice [ < size >gb ] : 5gb
1 out of Max. 16 slices created, 15 available.

# volslice create s0 -s 1024 -z 5GB v1
1 out of Max. 16 slices created, 14 available.
```

volslice remove

Synopsis

```
volslice remove slice_name
```

```
volslice remove -v volume_name
```

Function

Removes a volume slice. If a slice name is provided, the named slice is removed. If a volume name is provided, all slices of that volume are removed, with the underlying volume remaining intact.

This command also causes the slice information to be deleted from the disk and the dynamic data structure. All the Persistent Group Reservations (PGR) and LUN mapping and masking permissions are also deleted.

Parameters

TABLE 12-25 volslice remove Command Options and Parameters

Parameter	Function
<i>slice_name</i>	Specifies the slice name such as s0, s1, s2, and s3..
-v <i>volume_name</i>	Specifies the volume name such as v0, v1, v2, and v3 for which all slices are to be removed.

Example

- The first example removes slice name s4.

```
:/: volslice remove s4
WARNING - The slice will be deleted.
Continue ? [N]: Y
```

- The second example removes all slices from volume v1.

```
:/: volslice remove -v v1
WARNING - All slices in the given volume will be deleted.
Continue? [N]: Y
```

LUN Mapping Command Descriptions

This section describes the following commands:

- “lun map list” on page 166
- “lun map add” on page 168
- “lun map rm” on page 169
- “lun map rm all” on page 170

lun map list

Synopsis

```
lun map list
```

```
lun map list slice slice-no.
```

```
lun map list lun lun-no.
```

Function

Without any options, this command lists all current LUN and slice mapping entries. With an option, it displays the corresponding LUN/slice pair compliment for a specific slice number or for a LUN number.

Parameters

TABLE 12-26 lun map list Command Options and Parameters

Parameter	Function
<i>slice-no.</i>	Specifies the slice number such as 0, 1, 2, 3, ... 63.
<i>lun-no.</i>	Specifies a LUN number such as 0, 1, 2, 3, ... 255.

Example

- The first example lists all current LUN and slice mapping entries.

```
# lun map list
Lun No      Slice No
0           0
1           1
2           2
3           3
```

- The second example gives the LUN mapping to slice 4.

```
# lun map list slice 4
Lun No      Slice No
4           4
```

- The third example gives the slice mapping to LUN 3.

```
# lun map list lun 3
Lun No      Slice No
3           3
```

lun map add

Synopsis

`lun map add lun lun-no. slice slice-no.`

Function

Maps a LUN to a slice number. A slice must be mapped to a corresponding LUN before it can be accessed by any HBA.

Parameters

TABLE 12-27 `lun map add` Command Options and Parameters

Parameter	Function
<i>lun-no.</i>	Specifies a LUN number such as 0, 1, 2, 3, ... 255.
<i>slice-no.</i>	Specifies the slice number such as 0, 1, 2, 3, ... 63.

Example

The following example maps LUN 8 to slice 3. This assumes LUN 8 has not been previously assigned.

```
# lun map add lun 8 slice 3
```

lun map rm

Synopsis

```
lun map rm lun lun-no
```

```
lun map rm lun lun-no. slice slice-no.
```

Function

Removes the mapping between a specified LUN and its corresponding slice.

Parameters

TABLE 12-28 lun map rm Command Options and Parameters

Parameter	Function
<i>lun-no</i> .	Specifies a LUN number such as 0, 1, 2, 3, ... 255.
<i>slice-no</i> .	Specifies the slice number such as 0, 1, 2, 3, ... 63.

Example

- The first example removes the existing LUN mapping to LUN 8.

```
# lun map rm lun 8
```

- The second example removes LUN mapping from LUN 8 to slice 3, assuming LUN 8 is mapped to slice 3. If LUN 8 is not mapped to slice 3, an error message is displayed: “Lun map entry for lun 8 does NOT exist.”

```
# lun map rm lun 8 slice 3
```

lun map rm all

Synopsis

```
lun map rm all
```

Function

Removes all current LUN mappings.

Example

```
# lun map rm all
```

LUN Masking Commands Descriptions

This section describes the following commands:

- “lun perm list” on page 172
- “lun perm” on page 174
- “lun default” on page 176
- “lun wwn list” on page 177
- “lun wwn rm” on page 178
- “lun wwn rm all” on page 179
- “hwwn listgrp” on page 181
- “hwwn list” on page 180
- “hwwn add” on page 182
- “hwwn rm” on page 183
- “hwwn rmgrp” on page 184

lun perm list

Synopsis

`lun perm list`

`lun perm list lun lun-no.`

`lun perm list wwn wwn-no.`

`lun perm list grp group-name`

Function

Lists the accessibility of LUNs based on specified criteria.

Parameters

TABLE 12-29 lun perm list Command Options and Parameters

Parameter	Function
<i>lun-no.</i>	Specifies a LUN number such as 0, 1, 2, 3, ... 255.
<i>wwn-no.</i>	Specifies a world wide number. This is a unique 8-byte port WWN of an HBA Fibre Channel port. The WWN must be exactly 16 alphanumeric characters consisting of a-f, A-F, and 0-9.
<i>group-name</i>	Specifies a specific group name such as node1, node2, etc. The group name has a maximum of 16 alphanumeric characters consisting of a-z, A-Z, 0-9, and -.

Example

- The first example lists all permissions for all LUNs. For each LUN, the default access permissions are followed by any exceptions.

```
# lun perm list
Lun  Slice  WWN                Group Name  Group Perm  WWN Perm  Effective Perm
-----
0    0      default            --          --          rw        rw
0    0      20020678ff345678  G1         none       ro        ro
1    1      default            --          --          ro        ro
1    1      20020678ff345678  G1         ro         rw        rw
1    1      20020678ee345678  G1         ro         none      ro
1    1      20020678ee345678  G2         ro         none      ro
1    1      20020678ab345678  G2         ro         none      ro
-----
```

- The second example lists all WWNs that can access LUN 0.

```
# lun perm list lun 0
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	none	none
0	0	20020678ff345678	--	--	rw	rw
0	0	20020678ff345678	--	--	rw	rw
0	0	20020678ff345678	--	--	ro	ro
0	0	20020678ff345678	--	--	rw	rw

- The third example lists all LUN's accessibility for WWN 20020678ff345678.

```
# lun perm list wwn 20020678ff345678
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	20020678ff345678	--	--	rw	rw
1	1	20020678ff345678	--	--	rw	rw
2	2	20020678ff345678	--	--	rw	rw

This WWN is not defined in any group.

- The fourth example lists all the permissions for members of group G1.

```
# lun perm list grp G1
```

Lun	Slice	WWN	Group Name	Group Perm	WWN Perm	Effective Perm
0	0	default	--	--	rw	rw
0	0	20020678ff345678	G1	none	ro	ro
1	1	default	--	--	ro	ro
1	1	20020678ff345678	G1	ro	rw	rw
1	1	20020678ff345678	G1	ro	none	rw
2	2	default	--	--	rw	none

lun perm

Synopsis

```
lun perm lun lun-no. rw | ro | none wwn wwn-no.
lun perm lun lun-no. rw | ro | none all_wwn
lun perm lun lun-no. rw | ro | none grp group-name
lun perm all_lun rw | ro | none wwn wwn-no.
lun perm all_lun rw | ro | none all_wwn
lun perm all_lun rw | ro | none grp group-name
```

Function

By default, no LUN access is given to any HBA. This command is used to grant LUN access to:

- A single HBA (specified by its WWN)
- A group of HBAs
- All HBAs

Parameters

TABLE 12-30 lun perm Command Options and Parameters

Parameter	Function
<i>lun-no.</i>	Specifies a LUN number such as 0, 1, 2, 3, ... 255.
<i>all_lun</i>	Specifies all LUNs.
<i>rw</i>	Specifies read and write permission.
<i>ro</i>	Specifies read-only permission.
<i>none</i>	Specifies no access permission—the host cannot see the LUN.
<i>wwn-no.</i>	Specifies a world-wide number. This is a unique 8-byte port WWN of an HBA Fibre Channel port. The WWN must be exactly 16 alphanumeric characters consisting of a-f, A-F, and 0-9.
<i>all_wwn</i>	Specifies all world-wide numbers.
<i>group-name</i>	Specifies a specific group name such as node1, node2, etc. The group name has a maximum of 16 alphanumeric characters consisting of a-z, A-Z, 0-9, and -.

Example

- The first example sets up LUN 3 with read/write access to WWN 20020da445678901.

```
# lun perm lun 3 rw wwn 20020da445678901
```

- The second example sets up LUN 2 with no access to all WWNs.

```
# lun perm lun 2 none all_wwn
```

- The third example sets up all LUNs with read-only access to group node1.

```
# lun perm all_lun ro grp node1
```

lun default

Synopsis

```
lun default lun lun-no. rw
lun default lun lun-no. ro
lun default lun lun-no. none
lun default all_lun rw
lun default all_lun ro
lun default all_lun none
```

Function

Sets up the LUN default permission. The default permission for an un-registered LUN is “none.”

Parameters

TABLE 12-31 lun default Command Options and Parameters

Parameter	Function
<i>lun-no.</i>	Specifies a LUN number such as 0, 1, 2, 3, ... 255.

Example

- The first example sets up the read/write default permission for LUN 3.

```
# lun default lun 3 rw
```

- The second example sets up the read only default permission for all LUNs.

```
# lun default all_lun ro
```

lun wwn list

Synopsis

```
lun wwn list
```

Function

Lists all registered WWNs.

Example

```
# lun wwn list
List of all registered WWNs
-----
2002067890123456
2002078901234567
2002089012345678
2002090123456789
```

lun wwn rm

Synopsis

```
lun wwn rm wwn WWN-no.
```

Function

Removes the specified WWN from the system database.

Parameters

TABLE 12-32 lun wwn rm Command Options and Parameters

Parameter	Function
<i>wwn-no.</i>	Specifies a world-wide number. This is a unique 8-byte port WWN of an HBA Fibre Channel port. The WWN must be exactly 16 alphanumeric characters consisting of a-f, A-F, and 0-9.

Example

In this example, running the `lun wwn list` command shows the registered WWNs. The first WWN is then removed with the `lun wwn rm` command. The result is then verified.

```
# lun wwn list
List of all registered WWNs
-----
2002067890123456
2002078901234567
2002089012345678
2002090123456789

# lun wwn rm wwn 2002067890123456
2002067890123456 will get default access right.
Do you want to remove? (Y/N)[n]:Y

# lun wwn list
List of all registered WWNs
-----
2002078901234567
2002089012345678
2002090123456789
```

lun wwn rm all

Synopsis

```
lun wwn rm all
```

Function

Removes all registered WWNs from the system database. The default LUN access permission applies to all unregistered WWNs.

Example

In this example, running the `lun wwn list` command shows the registered WWNs. The WWNs are then removed with the `lun wwn rm all` command. The result is then verified.

```
# lun wwn list
List of all registered WWNs
-----
2002067890123456
2002078901234567
2002089012345678
2002090123456789

# lun wwn rm all
Are you sure? (Y/N)[n]:Y

# lun wwn list
List of all registered WWNs
-----
```

hwwn list

Synopsis

`hwwn list group-name`

Function

Lists all the WWN(s) in a specified group.

Parameters

TABLE 12-33 `hwwn list` Command Options and Parameters

Parameter	Function
<i>group-name</i>	Specifies a specific group name such as node1 or node2. The group name has a maximum of 16 alphanumeric characters consisting of a-z, A-Z, 0-9, and -.

Example

```
# hwwn list node1
WWN in this group node1
-----
20020678ef345678
20020321ff345678
-----
** Total 2 entries **
```

hwwn listgrp

Synopsis

`hwwn listgrp`

Function

Lists all the registered group names.

Note – If all WWNs have been removed from a group, that group is not displayed in the output of any LUN masking command except this one. An empty group is always displayed in the output of the `hwwn listgrp` command.

Example

```
# hwwn listgrp
Group Name
-----
node1
node2
node3
-----
** Total 3 entries **
```

hwwn add

Synopsis

hwwn add

hwwn add *group-name* **wwn** *wwn-no*.

Function

Adds a WWN to an existing group. If a specified group name does not exist, the firmware creates a new group with that name and adds the WWN to it.

Note – If you issue the `hwwn add` command without any parameters, you will be interactively prompted to specify the parameters.

The maximum number of groups that can exist is 128.

Parameters

TABLE 12-34 `hwwn add` Command Options and Parameters

Parameter	Function
<i>group-name</i>	Specifies a specific group name such as <code>node1</code> or <code>node2</code> . The group name has a maximum of 16 alphanumeric characters consisting of a-z, A-Z, 0-9, and <code>-</code> .
<i>wwn-no</i>	Specifies a world-wide number. This is a unique 8-byte port WWN of an HBA Fibre Channel port. The WWN must be exactly 16 alphanumeric characters consisting of a-f, A-F, and 0-9.

Example

This example shows the interactive and manual methods to add the WWN `20020678ef345678` to group “`node1`.”

```
# hwwn add
Please enter Group Name: node1
Please enter WWN: 20020678ef345678
Add more WWN? [n]: N

# hwwn add node1 wwn 20020678ef345678
```

hwwn rm

Synopsis

hwwn rm

hwwn rm *group-name* **wwn** *wwn-no.*

Function

Removes one or more WWN(s) from a specific group.

Note – If you issue the `hwwn rm` command without any parameters, you will be interactively prompted to specify the parameters.

Parameters

TABLE 12-35 `hwwn rm` Command Options and Parameters

Parameter	Function
<i>group-name</i>	Specifies a specific group name such as <code>node1</code> or <code>node2</code> . The group name has a maximum of 16 alphanumeric characters consisting of a–z, A–Z, 0–9, and <code>-</code> .
<i>wwn-no.</i>	Specifies a world-wide number. This is a unique 8-byte port WWN of an HBA Fibre Channel port. The WWN must be exactly 16 alphanumeric characters consisting of a–f, A–F, and 0–9.

Example

This example shows the interactive and manual methods to remove the `12345678ef345678` and `87654321ff345678` WWNs from group `node1`.

```
# hwwn rm
Please enter Group Name: node1
Please enter WWN: 20020678ef345678
More WWN [n]: Y
Please enter WWN: 20020321ff345678
More WWN [n]: N

# hwwn rm node1 wwn 20020678ef345678
# hwwn rm node1 wwn 20020321ff345678
```

hwwn rmgrp

Synopsis

`hwwn rmgrp group-name`

Function

Removes a specific group name and all its associated WWNs.

Parameters

TABLE 12-36 `hwwn rmgrp` Command Options and Parameters

Parameter	Function
<i>group-name</i>	Specifies a specific group name such as node1 or node2. The group name has a maximum of 16 alphanumeric characters consisting of a-z, A-Z, 0-9, and -.

Example

```
# hwwn rmgrp node2
```

Fibre Channel Topology Command

The command in this section applies only to controller firmware version 2.1 and later. This section describes the following commands:

- “sys fc_topology” on page 186

sys fc_topology

Synopsis

```
sys fc_topology
```

```
sys fc_topology auto
```

```
sys fc_topology loop
```

```
sys fc_topology fabric_p2p
```

Function

Without any parameter, this command displays the current negotiated topology. See TABLE 12-38 and TABLE 12-39.

With a parameter selected, this command sets the T3+ array topology firmware mode.

Parameters

TABLE 12-37 hwnn add Command Options and Parameters

Parameter	Function
auto	Enables the controller firmware to automatically determine and set the topology firmware mode.
loop	Sets the loop topology firmware mode.
fabric_p2p	Sets the point-to-point topology firmware mode.

Note – Normally the automatic mode is set. The other modes can be used for special situations or for troubleshooting.

Example

- The first example displays the current negotiated network topology.

```
# sys fc_topology
Host port u1p1 Current Negotiated Topology=Auto
Point to Point, Port ID=100300
```

- The second example sets up the T3+ array topology firmware mode to automatic sense and set. For an explanation of the various current negotiated topology results, see TABLE 12-38 and TABLE 12-39 which show the topology inquiry results of the `sys fc_topology` command with different types of switches and T3 array topology firmware mode settings.

```
# sys fc_topology fabric_p2p
```

TABLE 12-38 Topology Inquiry Results for SANbox Switch

T3 Topology Mode Setting	SANbox Switch Port	Current Negotiated Topology	
		2.0.1 Firmware Results	2.1 Firmware Results
auto	F-Port	-	Fabric Point-to-Point
loop	F-Port	Invalid	Invalid
fabric_p2p	F-Port	-	Fabric Point-to-Point
auto	TL-Port	-	Arbitrated loop
loop	TL-Port	Arbitrated loop	Arbitrated loop
fabric_p2p	TL-Port	-	Invalid
auto	SL-Port	-	Arbitrated loop
loop	SL-Port	Arbitrated loop	Arbitrated loop
fabric_p2p	SL-Port	-	Invalid
auto	Fabric	-	Fabric loop
loop	Fabric	Fabric loop	Fabric loop
fabric_p2p	Fabric	-	Fabric Point-to-Point

TABLE 12-39 Topology Inquiry Results for Brocade Switch

T3 Topology Mode Setting	Brocade Switch Port	Current Negotiated Topology	
		2.0.1 Firmware Results	2.1 Firmware Results
auto	qlPortEnable	-	Arbitrated loop
loop	qlPortEnable	Arbitrated loop	Arbitrated loop
fabric_p2p	qlPortEnable	-	Invalid

TABLE 12-39 Topology Inquiry Results for Brocade Switch *(Continued)*

T3 Topology Mode Setting	Brocade Switch Port	Current Negotiated Topology	
		2.0.1 Firmware Results	2.1 Firmware Results
auto	qlPortDisable	-	Fabric Point-to-Point
loop	qlPortDisable	Invalid	Invalid
fabric_p2p	qlPortDisable	-	Fabric Point-to-Point

Network Time Protocol Commands

The commands in this section apply only to controller firmware version 2.1 and later. This section describes the following commands:

- “set timezone” on page 190
- “ntp” on page 196
- “ntp poll” on page 198
- “ntp interval” on page 199
- “ntp server” on page 200
- “ntp stats” on page 197

set timezone

Synopsis

set timezone :*region*

Function

Sets the time zone of the Server that is running the NTP daemon. The time zone region codes are provided in TABLE 12-41.

Parameters

TABLE 12-40 set timezone Command Options and Parameters

Parameter	Function
<i>region</i>	Establishes the timezone for the system. See TABLE 12-41.

Example

```
# set timezone :America/Los_Angeles
```

TABLE 12-41 Olson Time Zone Regions

Africa/Abidjan	Africa/Accra	Africa/Addis_Ababa	Africa/Algiers
Africa/Asmera	Africa/Bamako	Africa/Bangui	Africa/Banjul
Africa/Bissau	Africa/Blantyre	Africa/Brazzaville	Africa/Bujumbura
Africa/Cairo	Africa/Casablanca	Africa/Ceuta	Africa/Conakry
Africa/Dakar	Africa/Dar_es_Salaam	Africa/Djibouti	Africa/Douala
Africa/El_Aaiun	Africa/Freetown	Africa/Gaborone	Africa/Harare
Africa/Johannesburg	Africa/Kampala	Africa/Khartoum	Africa/Kigali
Africa/Kinshasa	Africa/Lagos	Africa/Libreville	Africa/Lome
Africa/Luanda	Africa/Lubumbashi	Africa/Lusaka	Africa/Malabo
Africa/Maputo	Africa/Maseru	Africa/Mbabane	Africa/Mogadishu

TABLE 12-41 Olson Time Zone Regions (Continued)

Africa/Monrovia	Africa/Nairobi	Africa/Ndjamena	Africa/Niamey
Africa/Nouakchott	Africa/ Ouagadougou	Africa/Porto-Novo	Africa/Sao_Tome
Africa/Timbuktu	Africa/Tripoli	Africa/Tunis	Africa/Windhoek
America/Adak	America/ Anchorage	America/Anguilla	America/Antigua
America/ Araguaina	America/Aruba	America/Asuncion	America/Atka
America/Barbados	America/Belem	America/Belize	America/Boa_Vista
America/Bogota	America/Boise	America/ Buenos_Aires	America/ Cambridge_Bay
America/Cancun	America/Caracas	America/ Catamarca	America/Cayenne
America/Cayman	America/Chicago	America/ Chihuahua	America/Cordoba
America/ Costa_Rica	America/Cuiaba	America/Curacao	America/ Danmarkshavn
America/Dawson	America/ Dawson_Creek	America/Denver	America/Detroit
America/Dominica	America/Edmonton	America/Eirunepe	America/El_Salvador
America/Ensenada	America/ Fort_Wayne	America/Fortaleza	America/Glace_Bay
America/Godthab	America/ Goose_Bay	America/ Grand_Turk	America/Grenada
America/ Guadeloupe	America/ Guatemala	America/Guayaquil	America/Guyana
America/Halifax	America/Havana	America/ Hermosillo	America/Indiana
America/Indiana/ Indianapolis	America/Indiana/ Knox	America/Indiana/ Marengo	America/Indiana/ Vevay
America/ Indianapolis	America/Inuvik	America/Iqaluit	America/Jamaica
America/Jujuy	America/Juneau	America/Kentucky	America/Kentucky/ Louisville
America/Kentucky /Monticello	America/Knox_IN	America/La_Paz	America/Lima
America/ Los_Angeles	America/Louisville	America/Maceio	America/Managua

TABLE 12-41 Olson Time Zone Regions (Continued)

America/Manaus	America/ Martinique	America/Mazatlan	America/Mendoza
America/ Menominee	America/Merida	America/ Mexico_City	America/Miquelon
America/ Monterrey	America/ Montevideo	America/Montreal	America/Montserrat
America/Nassau	America/ New_York	America/Nipigon	America/Nome
America/Noronha	America/ North_Dakota	America/ North_Dakota/ Center	America/Panama
America/ Pangnirtung	America/Paramaribo	America/Phoenix	America/Port-au-Prince
America/Port_of_Spain	America/ Porto_Acre	America/ Porto_Velho	America/Puerto_Rico
America/ Rainy_River	America/ Rankin_Inlet	America/Recife	America/Regina
America/ Rio_Branco	America/Rosario	America/Santiago	America/ Santo_Domingo
America/Sao_Paulo	America/ Scoresbysund	America/Shiprock	America/St_Johns
America/St_Kitts	America/St_Lucia	America/ St_Thomas	America/St_Vincent
America/ Swift_Current	America/ Tegucigalpa	America/Thule	America/Thunder_Bay
America/Tijuana	America/Tortola	America/ Vancouver	America/Virgin
America/ Whitehorse	America/Winnipeg	America/Yakutat	America/Yellowknife
Antarctica	Antarctica/Casey	Antarctica/Davis	Antarctica/ DumontDURville
Antarctica/Mawson	Antarctica/ McMurdo	Antarctica/Palmer	Antarctica/South_Pole
Antarctica/Syowa	Antarctica/Vostok	Arctic/ Longyearbyen	Asia/Aden
Asia/Almaty	Asia/Amman	Asia/Anadyr	Asia/Aqtau
Asia/Aqtobe	Asia/Ashgabat	Asia/Ashkhabad	Asia/Baghdad
Asia/Bahrain	Asia/Baku	Asia/Bangkok	Asia/Beirut

TABLE 12-41 Olson Time Zone Regions (Continued)

Asia/Bishkek	Asia/Brunei	Asia/Calcutta	Asia/Choibalsan
Asia/Chongqing	Asia/Chungking	Asia/Colombo	Asia/Dacca
Asia/Damascus	Asia/Dhaka	Asia/Dili	Asia/Dubai
Asia/Dushanbe	Asia/Gaza	Asia/Harbin	Asia/Hong_Kong
Asia/Hovd	Asia/Irkutsk	Asia/Istanbul	Asia/Jakarta
Asia/Jayapura	Asia/Jerusalem	Asia/Kabul	Asia/Kamchatka
Asia/Karachi	Asia/Kashgar	Asia/Katmandu	Asia/Krasnoyarsk
Asia/ Kuala_Lumpur	Asia/Kuching	Asia/Kuwait	Asia/Macao
Asia/Magadan	Asia/Manila	Asia/Muscat	Asia/Nicosia
Asia/Novosibirsk	Asia/Omsk	Asia/Phnom_Penh	Asia/Pontianak
Asia/Pyongyang	Asia/Qatar	Asia/Rangoon	Asia/Riyadh
Asia/Riyadh87	Asia/Riyadh88	Asia/Riyadh89	Asia/Saigon
Asia/Sakhalin	Asia/Samarkand	Asia/Seoul	Asia/Shanghai
Asia/Singapore	Asia/Taipei	Asia/Tashkent	Asia/Tbilisi
Asia/Tehran	Asia/Tel_Aviv	Asia/Thimbu	Asia/Thimphu
Asia/Tokyo	Asia/ Ujung_Pandang	Asia/Ulaanbaatar	Asia/Ulan_Bator
Asia/Urumqi	Asia/Vientiane	Asia/Vladivostok	Asia/Yakutsk
Asia/Yekaterinburg	Asia/Yerevan	Atlantic/Azores	Atlantic/Bermuda
Atlantic/Canary	Atlantic/ Cape_Verde	Atlantic/Faeroe	Atlantic/Jan_Mayen
Atlantic/Madeira	Atlantic/Reykjavik	Atlantic/ South_Georgia	Atlantic/St_Helena
Atlantic/Stanley	Australia	Australia/ACT	Australia/Adelaide
Australia/Brisbane	Australia/ Broken_Hill	Australia/Canberra	Australia/Darwin
Australia/Hobart	Australia/LHI	Australia/ Lindeman	Australia/Lord_Howe
Australia/ Melbourne	Australia/NSW	Australia/North	Australia/Perth
Australia/ Queensland	Australia/South	Australia/Sydney	Australia/Tasmania
Australia/Victoria	Australia/West	Australia/ Yancowinna	Brazil/Acre

TABLE 12-41 Olson Time Zone Regions (Continued)

Brazil/DeNoronha	Brazil/East	Brazil/West	CET
CST6CDT	Canada/Atlantic	Canada/Central	Canada/East-Saskatchewan
Canada/Eastern	Canada/Mountain	Canada/Newfoundland	Canada/Pacific
Canada/Saskatchewan	Canada/Yukon	Chile/Continental	Chile/EasterIsland
Cuba	EET	EST	EST5EDT
Egypt	Eire	Etc/GMT	Etc/GMT+0
Etc/GMT+1	Etc/GMT+10	Etc/GMT+11	Etc/GMT+12
Etc/GMT+2	Etc/GMT+3	Etc/GMT+4	Etc/GMT+5
Etc/GMT+6	Etc/GMT+7	Etc/GMT+8	Etc/GMT+9
Etc/GMT-0	Etc/GMT-1	Etc/GMT-10	Etc/GMT-11
Etc/GMT-12	Etc/GMT-13	Etc/GMT-14	Etc/GMT-2
Etc/GMT-3	Etc/GMT-4	Etc/GMT-5	Etc/GMT-6
Etc/GMT-7	Etc/GMT-8	Etc/GMT-9	Etc/GMT0
Etc/Greenwich	Etc/UCT	Etc/UTC	Etc/Universal
Etc/Zulu	Europe/Amsterdam	Europe/Andorra	Europe/Athens
Europe/Belfast	Europe/Belgrade	Europe/Berlin	Europe/Bratislava
Europe/Brussels	Europe/Bucharest	Europe/Budapest	Europe/Chisinau
Europe/Copenhagen	Europe/Dublin	Europe/Gibraltar	Europe/Helsinki
Europe/Istanbul	Europe/Kaliningrad	Europe/Kiev	Europe/Lisbon
Europe/Ljubljana	Europe/London	Europe/Luxembourg	Europe/Madrid
Europe/Malta	Europe/Minsk	Europe/Monaco	Europe/Moscow
Europe/Nicosia	Europe/Oslo	Europe/Paris	Europe/Prague
Europe/Riga	Europe/Rome	Europe/Samara	Europe/San_Marino
Europe/Sarajevo	Europe/Simferopol	Europe/Skopje	Europe/Sofia
Europe/Stockholm	Europe/Tallinn	Europe/Tirane	Europe/Tiraspol
Europe/Uzhgorod	Europe/Vaduz	Europe/Vatican	Europe/Vienna
Europe/Vilnius	Europe/Warsaw	Europe/Zagreb	Europe/Zaporozhye

TABLE 12-41 Olson Time Zone Regions (Continued)

Europe/Zurich	GB	GB-Eire	GMT
GMT+0	GMT-0	GMT0	Greenwich
HST	Hongkong	Iceland	Indian/Antananarivo
Indian/Chagos	Indian/Christmas	Indian/Cocos	Indian/Comoro
Indian/Kerguelen	Indian/Mahe	Indian/Maldives	Indian/Mauritius
Indian/Mayotte	Indian/Reunion	Iran	Israel
Jamaica	Japan	Kwajalein	Libya
MET	MST	MST7MDT	Mexico/BajaNorte
Mexico/BajaSur	Mexico/General	Mideast/Riyadh87	Mideast/Riyadh88
Mideast/Riyadh89	NZ	NZ-CHAT	Navajo
PRC	PST8PDT	Pacific/Apia	Pacific/Auckland
Pacific/Chatham	Pacific/Easter	Pacific/Efate	Pacific/Enderbury
Pacific/Fakaofu	Pacific/Fiji	Pacific/Funafuti	Pacific/Galapagos
Pacific/Gambier	Pacific/Guadalcanal	Pacific/Guam	Pacific/Honolulu
Pacific/Johnston	Pacific/Kiritimati	Pacific/Kosrae	Pacific/Kwajalein
Pacific/Majuro	Pacific/Marquesas	Pacific/Midway	Pacific/Nauru
Pacific/Niue	Pacific/Norfolk	Pacific/Noumea	Pacific/Pago_Pago
Pacific/Palau	Pacific/Pitcairn	Pacific/Ponape	Pacific/Port_Moresby
Pacific/Rarotonga	Pacific/Saipan	Pacific/Samoa	Pacific/Tahiti
Pacific/Tarawa	Pacific/Tongatapu	Pacific/Truk	Pacific/Wake
Pacific/Wallis	Pacific/Yap	Poland	Portugal
ROC	ROK	Singapore	SystemV/AST4
SystemV/AST4ADT	SystemV/CST6	SystemV/CST6CDT	SystemV/EST5
SystemV/EST5EDT	SystemV/HST10	SystemV/MST7	SystemV/MST7MDT
SystemV/PST8	SystemV/PST8PDT	SystemV/YST9	SystemV/YST9YDT
Turkey	UCT	US/Alaska	US/Aleutian
US/Arizona	US/Central	US/East-Indiana	US/Eastern
US/Hawaii	US/Indiana-Starke	US/Michigan	US/Mountain
US/Pacific	US/Pacific-New	US/Samoa	UTC
Universal	W-SU	WET	Zulu

ntp

Synopsis

ntp

Function

Displays:

- The NTP server IP address
- The polling on/off status
- The polling interval in minutes.

Example

```
# ntp
server          22.3.36.42
poll            off
interval        5
```

ntp stats

Synopsis

`ntp stats`

Function

Displays:

- The time of the last poll
- The NTP server IP address
- The adjustment made to the T3 array internal clock
- The clock set success statement

Example

```
# ntp stats
lastpoll      Tue Feb 19 21:07:32 GMT 2002
server        22.3.36.42
offset        + 0.00413176
status        Successfully adjusted the time.
```

ntp poll

Synopsis

```
ntp poll unicast
```

```
ntp poll off
```

Function

Turns NTP polling on and off.

Parameters

TABLE 12-42 ntp poll Command Options and Parameters

Parameter	Function
unicast	Turns polling on.
off	Turns polling off.

Example

```
# ntp poll unicast
```

ntp interval

Synopsis

```
ntp interval time
```

Function

Sets the polling interval from 1 to 60 minutes. The default is 10 minutes.

Parameters

TABLE 12-43 `ntp interval` Command Options and Parameters

Parameter	Function
<i>time</i>	Sets polling interval time in minutes from 1 to 60.

Example

```
# ntp interval 5
```

ntp server

Synopsis

`ntp server IP_address`

Function

Establishes the server that is running the Network Time Protocol daemon in the T3 array firmware.

Parameters

TABLE 12-44 `ntp server` Command Options and Parameters

Parameter	Function
<i>IP_address</i>	Specifies the IP address of the server running the NTP daemon.

Example

```
# ntp server 22.3.36.42
```

syslog Error Messages

This chapter contains a list of errors that can be reported by the array, along with error definitions.

Topics in this chapter include:

- “Error Message Syntax” on page 202
- “Reset Log Types” on page 219
- “Command-Line Error Messages” on page 220
- “Examples” on page 230

Error Message Syntax

Error message syntax consists of the following three components:

- “Message Types” on page 202
- “FRU Identifiers” on page 203
- “Error Message Listing” on page 203

These components are described in the following subsections, which include a listing of `syslog` error and warning messages.

Message Types

A `syslog` daemon exists in the hardware RAID controller that records system messages and provides remote monitoring. The four possible levels of messages are listed in TABLE 13-1 in order of their severity.

TABLE 13-1 Message Types

Message Type	Definition
Error	Indicates a critical system event requiring immediate user intervention or attention. For example, an over-temperature condition exists or a detected FRU is being removed.
Warning	Indicates a possible event requiring eventual user intervention. For example, a FRU has been disabled.
Notice	Indicates a system event that may be a side effect of other events or may be a normal condition. For example, the power switch is turned off.
Information	Indicates a system event that has no consequence on the running health of the system. For example, notification of a of a proper functioning FRU.

FRU Identifiers

The syntax of the error message uses a FRU identifier to refer to a particular FRU in an array. This identifier contains a unit constant (*u*), the unit number (*n*), the FRU constant (*ctr* for controller card, *pcu* for power and cooling unit, *1* for interconnect card, *d* for disk drive), and the FRU number (*n*).

TABLE 13-2 FRU Identifiers

FRU	Identifier	Unit Number
Controller card	<i>uencidctr</i>	<i>encid</i> = unit number (1, 2)
Power and cooling unit	<i>uencidpcu</i> [1 2]	<i>encid</i> = unit number (1, 2) <i>n</i> = pcu number (1, 2)
Interconnect card	<i>uencid1</i> [1 2]	<i>encid</i> = unit number (1, 2) <i>n</i> = interconnect card number (1, 2)
Disk drive	<i>uencidn</i>	<i>encid</i> = unit number (1, 2) <i>n</i> = disk drive number (1 to 9)

Error Message Listing

TABLE 13-3 lists possible error messages.

TABLE 13-3 Error Messages

Error Message	Description	Solution
System		
E: No boot image on disk, writing EPROM to boot image area...	No master copy of image on disk. Writing to disk from flash memory.	
E: EPROM update failed...	Failed to write image from disk to flash memory.	
E: Power On Self Test failed...	This is a fatal error that means the firmware probably could not talk to the drives. This could have been caused by the interconnect cables, interconnect cards, the controller, disks/firmware, or a bad centerplane.	Troubleshoot the unit. Check its status.
Controller Card		

TABLE 13-3 Error Messages *(Continued)*

Error Message	Description	Solution
E: u<n>ctr: Not present	The array detects the controller is missing; the controller has been removed.	Reseat or replace the controller card.
E: u<n>ctr: Missing; system shutting down in minutes	A controller card is missing and a countdown has started. If the controller is not replaced in the time frame, then the system will shut down.	Replace the controller.
E: u<n>ctr: Missing; system shutting down now	A controller card was not replaced within 30-minute time frame; system shutting down now.	Replace the controller card and power on the system.
E: u<n>ctr BIST: ISP2200 test failed	ISP2200 POST test failed.	Please contact your contract service provider for further analysis.
E: u<n>ctr BIST: Data cache memory test failed	Cache Memory POST test failed.	Please contact your contract service provider for further analysis.
E: u<n>ctr BIST: XOR functions and datapaths test failed	XOR Memory POST test failed. Controller will be reset.	Controller board replacement may be necessary. Please take appropriate action.

TABLE 13-3 Error Messages (Continued)

Error Message	Description	Solution
<p>E: u<n>ctr XOR: Flags=<flags> Cntr=<error count> Synd= syndrome> Addr=<addr></p>	<p>An ECC soft/hard error is detected and the error is a multi-bit uncorrectable error.</p> <p><flags> - possible values are:</p> <ul style="list-style-type: none"> • A - Accumulator VRAM error • B - Buffer (system cache SDRAM) DIMM error • S - Single-bit error • M - Multi-bit error • V - Accumulator drive out error • D - Buffer drive out error • c - Control signal drive out error • a - Address line drive out error <p><error count> - ECC error count after last report</p> <p><syndrome> - ECC error syndrome</p> <p><addr> - Offending address</p>	<p>If problem persists, please contact your Sun authorized service provider.</p>
<p>E: XOR: Acc VRAM address line drive out error</p>	<p>An ECC uncorrectable error occurred on the address output line of XOR FPGA. This indicates a stuck-a-0/stuck-a-1, or shorted address line.</p>	<p>Please contact your Sun service provider.</p>
<p>E: XOR: Acc VRAM control line drive out error</p>	<p>An ECC uncorrectable error occurred on the control output line of XOR FPGA. This indicates a stuck-a-0/stuck-a-1, or shorted control line.</p>	<p>Please contact your Sun service provider.</p>
<p>E: XOR: Buf mem SDRAM address line drive out error</p>	<p>An ECC uncorrectable error occurred on the address output line of XOR FPGA. This indicates a stuck-a-0/stuck-a-1, or shorted address line.</p>	<p>Please contact your Sun service provider.</p>
<p>E: XOR: Buf mem SDRAM control line drive out error</p>	<p>An ECC uncorrectable error occurred on the control output line of XOR FPGA. This indicates a stuck-a-0/stuck-a-1, or shorted control line.</p>	<p>Please contact your Sun service provider.</p>

TABLE 13-3 Error Messages (Continued)

Error Message	Description	Solution
<p>E: u<n>ctr Diag Fail: Test System Memory failed or E: u<n>ctr Diag Fail: Test ISP2200 failed or E: u<n>ctr Diag Fail: Test data cach memory failed or E: u<n>ctr Diag Fail: XOR Memory Test failed. System Resetting...</p>		<p>Results of diagnostic tests, should be preceded by more detailed error messages.</p>
Interconnect (loop) Card and Cable		
<p>E: u<n>l<n>: Missing; system shutting down in minutes</p>	<p>A loopcard is missing, and a countdown has started. If the loop card is not replaced in the time frame, then the system will shut down.</p>	<p>Replace the loop card.</p>
<p>E: u<n>l<n>: Missing; system shutting down now</p>	<p>A loopcard was not replaced within 30-minute time frame; system shutting down now.</p>	<p>Replace the loop card and power on the system.</p>
<p>E: u<n>l<n>: Not present</p>	<p>System detects the interconnect card is missing.</p>	<p>Reseat or replace the interconnect card.</p>
<p>E: u<n>l<n>: UP cable not present or E: u<n>l<n>: DOWN cable not present</p>	<p>System detects missing interconnect cable for interconnect card.</p>	<p>Check cables.</p>
Disk Drive		
<p>E: u<n>d<n>: Missing; system shutting down in minutes</p>	<p>A disk drive is missing, and a countdown has begun. If the drive is not replaced in the time frame, then the system will shut down.</p>	<p>Replace the missing disk drive.</p>
<p>E: u<n>d<n>: Missing; system shutting down now</p>	<p>A disk drive was not replaced within the 30-minute time frame; system shutting down now.</p>	<p>Replace the disk drive and power on the system.</p>

TABLE 13-3 Error Messages (Continued)

Error Message	Description	Solution
E: Multi-disk failure, access disallowed	System detects multiple disk failure in the LUN, and automatically unmounts this volume.	Check the syslog for any other related messages. Replace the failed disks.
E: u<n>d<n>: Not present	System detects a drive was removed.	Reseat or replace the disk drive.
E: FATAL: No disks were found during boot up or E: FATAL: Check hardware run diags or E: FATAL: Fix problem and try rebooting	During boot up, no disks were found and the system will reboot.	Check the disks and run diagnostics.
Power and Cooling Unit (PCU)		
E: u<n>pcu<n>: Missing; system shutting down in minutes	A PCU is missing and a countdown has started. If the PCU is not replaced in the time frame, the system will shut down.	Replace the PCU.
E: u<n>pcu<n>: Missing; system shutting down now	A PCU was not replaced within the 30-minute time frame; system shutting down now.	Replace the PCU; power on the system.
E: u<n>pcu<n> Not present	System detects the PCU is missing.	Reseat or replace the PCU.
E: u<n>pcu<n>: Over temperature, serial no = x	System detects that a PCU has reached an over temperature state.	Replace the power and cooling unit.
E: u<n>pcu<n>: Battery not present	System detects that the battery is missing.	Replace the PCU.
E: u<n>ctrl: Multiple Fan Faults; system shutting down in minutes	The controller has detected a fault on fans in both PCUs.	Determine which fans in which PCU failed; replace the PCU.
E: u<n>ctrl: Multiple Fan Faults; system shutting down now	The controller has detected a fault on fans in both PCUs.	Determine which fans in which PCU failed; replace the PCU.

Warning Message Listing

TABLE 13-4 lists possible warning messages, which are divided into FRU categories and System.

TABLE 13-4 Warning Messages

Warning Message	Description	Solution
System		
W: u<n>ctr recon failed in vol volume_name	Reconstruct operation failed for this volume.	Verify that the volume is still mounted and that there is only one drive disabled. Restart the reconstruct operation manually, or reconstruct the data to another drive. Previous messages in <code>syslog</code> should indicate which disk(s) had a problem.
W: u<n>ctr initialization failed in vol volume_name	Volume data initialization to zero failed.	Check the disk drive status using the <code>vol stat</code> command. If any disk has a nonzero status, either initialize the volume using the <code>vol init</code> command, or replace the disk.

TABLE 13-4 Warning Messages *(Continued)*

Warning Message	Description	Solution
W: u<n>ctr Reset (xxxx) occurred or	A hard reset occurred. Message indicates a controller has been reset manually.	Check adjacent messages to see if this event indicates an error.
W: u<n>ctr Hardware Reset (xxxx) occurred or		
W: u<n>ctr Exception Reset (xxxx) occurred or		
W: u<n>ctr Assertion Reset (xxxx) occurred or		
W: u<n>ctr RaidFail Reset (xxxx) occurred or		
W: u<n>ctr Takeover Reset (xxxx) occurred or		
W: u<n>ctr PsosFail Reset (xxxx) occurred or		
W: u<n>ctr SysFail Reset (xxxx) occurred or		
W: u<n>ctr MmuFail Reset (xxxx) occurred		

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
<p>W: u<n>ctr Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr Hardware Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr Exception Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr Assertion Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr RaidFail Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr Takeover Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr PsosFail Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr SysFail Reset (xxxx) was initiated at <date> <time> <message></p> <p>or</p> <p>W: u<n>ctr MmuFail Reset (xxxx) was initiated at <date> <time> <message></p>	<p>A soft reset occurred, for one of several reasons. The specific error is indicated by the type.</p>	<p>Check Table B-5 for information about the type of reset that occurred. The field here should match one of those listed in Table B-5.</p>
<p>W: u<n>ctr Illegal mp_support mode for explicit LUN fail over request (opcode 0xD0 received) on LUN</p>	<p>The mp_support option is set to an invalid value.</p>	<p>Use the sys list command (or Component Manager) to set the mp_support flag to mpzio or std.</p>

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
<pre>W: u<n>d<n>: SVD_PATH_FAILOVER: path_id = x or W: u<n>d<n>: SVD_PATH_FAILOVER: path_id = x, lid = x</pre>	<p>A failover has occurred, indicating that one of the system's replaceable units is faulty.</p>	<p>Troubleshoot and replace the failed component. You can find more information in related syslog messages and by using the <code>fru stat</code> command.</p>
<pre>W: u<n>l<n> LMON predictive failure threshold exceeded - LIP</pre>	<p>The Fibre Channel loop is exhibiting an unexpected number of LIPs. This problem can be generated from any of the FRUs.</p>	<p>In certain cases, running the <code>ofdg</code> diagnostic may find the problem. Notify your Sun authorized service provider.</p>
<pre>W: Could not retrieve header information of level N</pre>	<p>When the <code>id read</code> or <code>fru stat</code> commands fail to retrieve header information from the flash PROM, Level 2 or Level 3 code may be corrupted.</p>	<p>Try the command(s) again. If problem persists, please contact your Sun authorized service provider.</p>
<pre>W: u<n>l<n> cache memory size is less than 256 Mbytes</pre>	<p>The cache memory is smaller than 256 Mbytes. The T3+ expects a minimum of 256 Mbytes of cache.</p>	
Controller Card		
<pre>W: u<n>ctr: Offline</pre>	<p>A controller has stopped communicating with the rest of the system. Possibly the connection between the controller and the rest of the system has been disrupted.</p>	<p>Check that the controller card is connected securely. If its state becomes disabled, replace the controller card.</p>
<pre>W: u<n>ctr: Inactive</pre>	<p>A controller is not communicating with the rest of the system. The controller will be disabled if it does not start communicating with the rest of the system within a certain time period.</p>	
<pre>W: u<n>ctr: Rebooting</pre>	<p>A controller has been rebooted because it was removed, disabled, or was offline and is now re-enabled. Alternatively, the system was shut down and restarted.</p>	<p>Check previous messages to determine if the reboot is an unexpected event.</p>

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>ctr: Failed to boot	The system was not able to reboot a controller successfully, even after repeated attempts.	Replace the controller.
W: u<n>ctr: Disabled	System disabled a controller.	Check the previous message queue to determine why the controller is disabled. Then enable or replace the controller.
W: u<n>ctr starting lun <lun_#> failover	A LUN failover has occurred.	Use the <code>sys stat</code> command to see if the controller is disabled. Then check the <code>syslog</code> for possible interconnect, host adapter, or cable problems.
W: u<n>ctr BIST: ECC mechanism verification failed		Contact your contract your Sun service provider for further analysis.
W: u<n>ctr: ISP2x00[N] qlcf_i_read_response: Debug Code - Immediate Notify CCBs Exhausted or W: u<n>ctr: ISP2x00[N] DMA Error Detected or W: u<n>ctr: ISP2x00[N] Received LOOP DOWN async event or W: u<n>ctr: ISP2x00[N] Bad request pkt or W: u<n>ctr: ISP2x00[N] Bad request pkt header	Indicates an error in one of the ISP 2x00 chips in the controller.	Notify your Sun authorized syour Sun ervice provider.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>ctr XOR T<n> ERROR: ACC should not be zero or W: u<n>ctr XOR T<n> ERROR: ACC should be zero or W: u<n>ctr XOR T<n> ERROR: BUF data not equal to previously written or W: u<n>ctr XOR T<n> ERROR: BUF data should be all zeros or W: u<n>ctr XOR T<n> ERROR: ORed data should be all zeros or W: u<n>ctr XOR T<n> ERROR: Compare error of BUS and orig data or W: u<n>ctr XOR T<n> ERROR: buf data should be all zeros or W: u<n>ctr XOR T<n> ERROR: Buf check finds all zeros or W: u<n>ctr XOR T<n> ERROR: Write to data cache mem failed or W: u<n>ctr XOR T<n> ERROR: Read from data cache mem failed or W: u<n>ctr XOR T<n> ERROR: Diag Fail	A XOR Memory POST test failed. The controller will be reset.	The controller board might need to be replaced. Please take appropriate action.
W: u<n>ctr ISP2x00 [N] LUN lun - INVALID command count	The Fibre-Channel driver has gone into an unknown state. This message might indicate a bug in the resource management logic.	Please contact your Sun authorized service provider.
W: u<n>ctr ISP2x00 [N] LUN lun Out of command resources	The fibre-channel driver resource management might need some tuning.	Please contact your Sun authorized service provider.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>ctr ISP2x00 [N] cmd = SCSI Command - REQUEST QUEUE FULL	The Fibre-Channel driver is not able to accept/send more commands from ISP. The driver retries automatically.	If this happens frequently, please contact your Sun authorized service provider.
W: u<n>ctr XOR: Flags=<flags> Cntr=<error count> Synd= syndrome> Addr=<addr>	An ECC soft/hard error is detected and the error is a single-bit correctable error. <flags> - possible values are: <ul style="list-style-type: none"> • A - Accumulator VRAM error • B - Buffer (system cache SDRAM) DIMM error • S - Single bit error • M - Multi-bit error • V - Accumulator drive out error • D - Buffer drive out error • c - Control signal drive out error • a - Address line drive out error <error count> - ECC error count after last report <syndrome> - ECC error syndrome <addr> - Offending address	The ECC error has been corrected.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
Interconnect (loop) card and Cable		
W: u<n>l<n>: Offline	The system is not able to poll the interconnect card for information. If an interconnect card is disabled, it may cause other interconnect cards to go offline.	<ul style="list-style-type: none"> • If no interconnect card is disabled, check cable connections to the interconnect card and make sure it is firmly seated. • If the interconnect card does not come back online, replace the interconnect card.
W: u<n>l<n>: Disabled	The system detects that the interconnect card is disabled.	Check the previous message queue to determine why the interconnect card is disabled. Replace the interconnect card if necessary.
W: Loop 1 may be bad - Please run ofdg or W: Loop 2 may be bad - Please run ofdg	The system detects the possibility of an internal error in the loop interconnection.	Use the ofdg command to run the system's built-in loop diagnostic.
W: u<n>l<n> LMN2 predictive failure threshold exceeded - IO TIMEOUT	A timeout has occurred during the ondg test.	
W: u<n>l<n> Loop 1 may be bad - check interconnect cable or loopcard or W: u<n>l<n> Loop 2 may be bad - check interconnect cable or loopcard.	The system detects a potential loop failure.	<ol style="list-style-type: none"> 1. Check the syslog for any other related messages. 2. Use the fru stat command to determine the problem cause. 3. Remove and replace the bad part.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
Disk Drive		
W: u<n>d<n> could not open plugged disk	An attempt to open the newly plugged disk failed.	Unseat the drive and wait a few seconds before reseating it so that the system can retry this operation.
W: u<n>d<n> could not create system area	An attempt to a create system area on this disk failed.	<ol style="list-style-type: none"> 1. Unseat the drive and wait a few seconds before reseating it so that the system can retry this operation. 2. If step 1 fails, replace this drive with a new drive.
W: u<n>d<n> system area is bad	The system area of the drive has been corrupted.	<ol style="list-style-type: none"> 1. Unseat the drive and wait a few seconds before reseating it so that the system can retry this operation. 2. If step 1 fails, replace this drive with a new drive.
W: u<n>d<n> could not open disk, try unplug then plug	An attempt to bring newly installed drive online failed.	Unseat the drive and wait a few seconds before reseating it.
W: u<n>d<n> could not open disk to be checked	An attempt to open this drive failed.	<ol style="list-style-type: none"> 1. Unseat the drive, wait a few seconds, reseat it and retry the volume add procedure. 2. If step 1 fails, replace the drive with a new one.
W: u<n>d<n> system area verify failed	An attempt to verify the data in the system area failed.	<p>Replace the drive.</p> <ul style="list-style-type: none"> • If the volume exists, initialize the system area to fix the problem. • If the volume does not exist, retry the volume add procedure.
W: u<n>d<n> write disk failed, err=<errno>	An attempt to write this disk failed. The system will automatically disable this disk.	Check the <code>syslog</code> for any other related messages. Replace the failed disk.
W: u<n>d<n> Recon attempt failed	Disk Reconstruct failed for disk u<n>d<n>.	Check the <code>syslog</code> for other related failures. If the disk under reconstruction failed, replace the disk.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>d<n> Disable attempt failed	Disk Disable failed for disk u<n>d<n>.	Check if the disk is already disabled by using the <code>vol stat</code> command. If not disabled, disable it by typing the <code>vol disable u<n>d<n></code> command.
W: u<n>d<n> Installing u<n>d<n> failed, Try unplugging and then plugging	Disk drive hot plug installation failed.	Unseat the drive and reseal it.
W: u<n>d<n> Disk Bypassed	A drive is faulty and being bypassed from the data loop.	
W: u<n>ctr read failed during recon stripe scb=<n>	During reconstruction, another drive failed in the volume.	Check <code>syslog</code> for a multiple drive failure. Replace all failed disk drives.
W: u<n>d<n> hard err in vol starting auto disable	A hard error on this disk prompted a disable operation on this disk. The system will auto disable and reconstruct to the hot spare (standby), if there is one configured.	If disabled, replace the failed disk as soon as possible.
W: u<n>ctr disk error during recon, terminating recon	More than one drive in the volume might have failed.	Check the <code>syslog</code> file. Replace all failed disk drives.
W: u<n>d<n> SCSI Disk Error Occurred (path = 0x<n>, port = 0x<n>, lun = 0x<n>)	The SCSI disk drive returned an error status with the appropriate SCSI Sense Data.	Notify your Sun authorized service provider.
W: u<n>d<n> SCSI Disk Error Occurred (path = 0x<n>)	If the port has a value from 0 to 7, the message indicates a transient controller condition, not an error affecting a disk.	
W: Sense Key = 0x%x, Asc = 0x%x, Ascq = 0xN	The SCSI disk drive returned an error status with the appropriate SCSI Sense Data.	Notify your Sun authorized service provider.
W: Sense Data Description = xxxxxx		
W: Valid Information = 0xN		

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>d<n> is in wrong disk position	A drive is in the wrong position in the unit.	1. Try to determine the correct position for the disk from the syslog.
W: u<n>d<n> is in wrong disk position, previously in u<n>d<n>		2. Power off the system. 3. Reorder the affected disks. 4. Reboot the system.
W: Disk <n> may be bad - too slow	If the ondg mode is passive, drive <i>N</i> is bad.	Replace the drive.
W: Can't Disable Disk u<n>d<n> - xxx	When ondg detects a slowly degrading drive, if the ondg mode is active, the array will try to disable that drive. If the array can't disable the drive, it will print out this warning message with the error string.	
W: Disable Disk u<n>d<n>	Disk <i>d</i> is disabled.	
W: u<n>d<n> TMON over temperature threshold exceeded	A disk drive temperature threshold has been exceeded. The drive is spinning down.	Check the disk drives and possible airflow obstruction.
W: u<n>d<n> TMON warning temperature threshold exceeded	A disk drive temperature threshold has been exceeded; the drive is still online.	Check the disk drives and possible airflow obstruction.
W: u<n>d<n>: Failed	The system detects a disk failure and has automatically disabled the disk drive.	Check the syslog for any other related messages. Replace the failed disk.
Power and Cooling Unit (PCU)		
W: u<n>pcu<n>: Fan 1 failed, serial no = x	The system detects a fan fault on a PCU.	1. Verify that the batteries have had ample time to charge.
W: u<n>pcu<n>: Fan 2 failed, serial no = x		2. Ensure a refresh operation is not in progress. 3. Replace the PCU.
W: u<n>pcu<n>: DC not OK, serial no = x	System detects that the DC power source in the PCU is not functioning properly.	Replace the PCU.
W: u<n>pcu<n>: Disabled, serial no = x	A PCU has been disabled.	Check previous messages to see why the PCU has been disabled, and replace the PCU, if necessary.

TABLE 13-4 Warning Messages (Continued)

Warning Message	Description	Solution
W: u<n>pcu<n>: Off, serial no = x	A PCU is off.	Check <code>syslog</code> for a fan fault or PCU overtemp. Replace the PCU if necessary.
W: u<n>pcu<n>: On battery, serial no =x	The system detects that a PCU has switched to battery.	1. Make sure that the AC power cord is attached. 2. If the unit is receiving AC power, then replace the PCU.
W: u<n>pcu<n>: Switch off	The system detects a PCU has been powered off.	Power on the unit by pressing the power switch.
W: u<n>pcu<n>: Replace battery, hold time low, serial no = x	The system detects that a battery is nearing its end of life.	Replace the PCU.
W: u<n>pcu<n>: Battery life expired, replace battery	The system detects an expired battery.	Replace the PCU.
W: u<n>pcu<n>: Replace battery, battery has been on the shelf for more than two years	The battery's shelf life has exceeded its warranty period.	Replace the PCU.
W: u<n>pcu<n>: n days battery life left, replace battery	The system detects that battery is approaching its expiration date.	Replace the PCU.

Reset Log Types

If the error level is set to Notice (2) or higher (Warning or Error), you can trace the reason for the reset through the `syslog` file.

TABLE 13-5 Reset Log Types

Index	Type	Type Value	Description
0	Hardware	0x1000	User reset
1	Exception	0x2000	Exception
2	Assertion	0x3000	Software assertion
3	RaidFail	0x4000	RAID fatal error

TABLE 13-5 Reset Log Types (Continued)

Index	Type	Type Value	Description
4	Takeover	0x5000	Takeover
5	pSOS fail	0x6000	pSOS fatal error
6	SysFail	0x7000	System error

Command-Line Error Messages

The Sun StorEdge T3+ array issues a variety of error messages to the command line, indicating an incorrectly entered command or invalid operation. When you enter a command by itself, or the error is syntactic (for example, missing an argument or using the wrong format), the array displays the command synopsis. Otherwise, the array displays an error message consisting of a name in capital letters, a numerical code in hexadecimal digits, and a text message.

The following tables list the errors that the array displays. TABLE 13-6 lists the types of errors and the numerical range associated with each.

TABLE 13-6 CLI Error Message Types

Type of Error	Numerical Range	Description
Logical volume manager (LVM) driver error codes	0x10001–0x1000A	Disk drive-related errors
Virtual Node ("VN") error codes	0x200000–0x200025	Errors related to <code>vol</code> and other command-line operations
Port error codes	0x300000–0x300006	Errors related to the <code>port</code> command
Sys error codes	0x400000	Only one error, indicating an incorrect value
FRU error codes	0x500001–0x500076	Errors related to field replaceable units (FRUs)
pSOS operating system errors	00000001–C000FFFF	pSOS errors (embedded operating system)

RAID Errors and Other Common Errors

Volume-related errors (VN_ERRORS) are the most common error messages displayed. TABLE 13-7 lists the names and values of these errors.

TABLE 13-7 Volume-Related (VN) Errors

Error Name	Numerical Value	Message
VN_BADUNIT	0x200000	Bad unit number
VN_BADDRIVE	0x200001	Bad drive number
VN_BADPART	0x200002	Bad partition ID
VN_VOLEXISTS	0x200003	Volume already in use
VN_VOLNOTFOUND	0x200004	Volume name not found
VN_PARTHASFS	0x200005	Partition already has file system
VN_FACLOCKED	0x200006	Facility locked by other command
VN_BADATTR	0x200007	Unable to read attributes
VN_MOUNTED	0x200008	Volume already mounted
VN_UNMOUNTED	0x200009	Volume not mounted
VN_MNTINUSE	0x20000A	Mount point in use
VN_NOMEMORY	0x20000B	Could not allocate memory for operation
VN_ALREADYDSBL	0x20000C	Is already a disabled drive
VN_NODSBL	0x20000D	No drives are disabled
VN_ABORTED	0x20000E	Operation aborted
VN_NOTSUP	0x20000F	Operation not supported
VN_UNKVOL	0x200010	Unknown volume
VN_RAIDERR	0x200015	RAID error
VN_NOPART	0x200016	Partition has size 0
VN_PARTSMALL	0x200017	Partition too small
VN_UNKVIF	0x200019	Unknown interface
VN_UNKVIFTYP	0x20001A	Unknown interface type
VN_BADVOLNAME	0x20001B	Bad volume name
VN_BADVOLNAMELEN	0x20001C	Bad volume name too long

TABLE 13-7 Volume-Related (VN) Errors *(Continued)*

Error Name	Numerical Value	Message
VN_CFGNOTSUPPORTED	0x20001D	Unsupported volume configuration
VN_BADSTANDBYUNIT	0x20001E	Standby unit number is wrong
VN_DEVINVALID	0x20001F	Invalid drive specified
VN_LOCVOLBAD	0x200020	Local volume bad
VN_PORTMAPRM	0x200021	Volume still mapped to a port
VN_UNINITIALIZED	0x200022	Volume is uninitialized
VN_PENDING	0x200023	Operation is pending
VN_BADMODE	0x200024	Cache mode must be set to auto for mirroring
VN_MIRRORON	0x200025	Cannot change cache mode when mirroring is on
VN_CANTMOUNT	0x200026	Cannot mount volume because multiple disabled drives
VN_STARTPARAM	0x200027	Slice start parameter invalid or in use
VN_VSLBADNAME	0x200028	Bad slice name
VN_MAXSLICEERR	0x200029	No more slices can be created
VN_VSLNOTFOUND	0x20002A	Slice not found
VN_SIZEPARAM	0x20002B	Incorrect slice size parameter encountered
VN_VSLBADNAMELEN	0x20002C	Slice name exceeded 12 characters allowed
VN_VSLEXISTS	0x20002D	Slice name already exists
VN_NOSLICEINVOL	0x20002E	Volume does not have slice(s) to be removed
VN_VSLRAIDERR	0x20002F	RAID error in volume slicing

Some of these errors are generated more often than others, such as VN_MOUNTED, VN_UNMOUNTED, VN_MNTINUSE, VN_CFGNOTSUPPORTED, VN_DEVINVALID, VN_LOCVOLBAD, VN_UNINITIALIZED, VN_BADMODE, and VN_MIRRORON. In particular, VN_RAIDERR, code 0x200015, can result from a variety of scenarios and this warrants attention. The Sun StorEdge T3+ array uses a specific protocol for carrying out commands, and this protocol uses the RAID error as a general error message to be sent to the user. As such, the RAIDERR may result from any of a number of software or hardware conditions. Some cases involve problems related to the user

configuration, which can be remedied easily. Other cases are more subtle and are related to the function of the array's internal software. Detailed information for specific instances of the RAID error can be found in the `syslog`; an overview of the scenarios are provided here.

Categories of RAID Errors

The following is a list of the categories that the array's embedded protocol uses for RAID errors and a few of the cases within each category. Each category code, which can be useful in reference to the `syslog`, is included. Though not exhaustive, this list provides a general breakdown of common RAID error generators:

1. **Command Incomplete (0x1A):** The command was not executed correctly internally. The software responded to the command with too much or too little information. In certain cases, the command might simply be paused and will resume.
2. **Partial (conditional) Success (0x19):** This category includes the following cases:
 - a. **Aborting a nonexistent command:** A user has issued a command, then tried to abort it after the command is executed.
 - b. **Retry error:** The command was retried one or more times.
 - c. **Target error:** A volume is offline or disabled.
3. **Invalid Response (part of the category above; 0x19):** The software did not provide a valid response to the user command. These cases are more specific than the Command Incomplete category.
 - a. **Invalid information type (parameter):** The software responded with the wrong type of information.
 - b. **Error in information returned:** The information returned in response to the command is erroneous. This case implies an embedded software error.
 - c. **Function failed:** The command failed to retrieve the appropriate information.
 - d. **Zero size:** The command accessed a volume with zero size.
4. **Command Aborted(0x18):** The command is aborted, often because of a time-out. A command aborts when a component in the system freezes or if a connection is defective.
5. **Command Exception (0x17):** This category includes cases where the command cannot be executed. This error type is generated when a disabled, unavailable, or invalid drive or volume is specified. For example, after using a hot spare (standby) to reconstruct the data on a drive, you cannot refer to it again as a hot spare.

- a. Invalid name/address: Either the user or the internal software used a volume or disk drive name that is invalid or does not match the current configuration.
 - b. Invalid command fields: The command is no longer supported or the internal software used a command opcode that is unsupported.
 - c. Missing fields: The user or the internal software issued a command with missing information.
 - d. Drive (module) errors: The disk drive referenced may be unattached, disabled, substituted, or in the process of being reconstructed.
6. Machine Exception (0x16): This category includes cases where there is a hardware error, or where other commands are being executed, thereby giving a busy response.
- a. Drive fatal error: There is an error internal to a referenced drive.
 - b. Auto reconstruct or disable attempted: A drive that is being reconstructed or disabled is specified.
 - c. Queue full or busy response: The command cannot be executed because the system is busy processing other commands.
 - d. Unknown host: The specified host address is invalid or unreachable.
 - e. Single drive errors: A drive referenced by the command was not detected, the connection could not be opened, or the sysarea on the drive could not be created. This case implies that the drive or the connection to it is faulty. Alternatively, a retry of a command accessing a disk could fail to execute.
 - f. Multiple disk failure: An error occurred involving more than one drive.
 - g. Standby already in use: (This error resembles the one in the Command Exception category). In this case, the drive is busy processing a previous command. This case applies when the command is complete and if the drive configuration has changed as a result.
 - h. Volume (LUN) errors: A volume may be inaccessible, or its configuration may be corrupted and represented as nonvalid.
7. Intervention Required (0x14): Here, an error results when a volume is mounted or unmounted, in contrast to what is expected. Alternatively, a physical connection might be broken and should be reinstated (by replacing the appropriate FRUs).

RAID errors can result from invalid command arguments or from a system problem. The error might refer to the configuration of a volume or an individual drive. For example, you can encounter the error while reconfiguring the array with volumes that have been added but not mounted, or the problem might be related to the hardware or an embedded component.

In general, RAID errors can be diagnosed by checking the status of the volumes mounted on the array. Often, an existing but unmounted volume can trigger the error. Other times, a conflict occurs when a new version of the binary is downloaded while previous volumes are still being used.

Investigating RAID Errors

Some guidelines for investigating RAID errors include:

- 1. Check the state of the current volumes with the `vol stat` command.**
 - If the volumes are unmounted, try remounting them and then resetting the system using the `reset` command.
 - If you are unable to remount the volumes, try removing all the volumes, resetting the system, then adding them back before remounting them.
- 2. Check the host connection to the array.**

On hosts running the Solaris software environment, the `format` command should match the number of volumes present on the array. If the number of volumes listed do not match, see Chapter 5 for troubleshooting instructions. Specifically, the T300 entries listed by the `format` command should be recognized and labeled, and the number of these entries should equal the number of volumes mounted on the array.

- 3. If you suspect that a hardware problem is causing the RAID error, use the `fru list` and `fru stat` commands to check the status of the components.**

It might also be worthwhile to check the cables and connections between partner-group units and between the host and the array units.

More detail regarding the error might be available in the `syslog`, in which case note the date and time of the error for lookup. However, most common cases can be handled as described above.

Port Errors

TABLE 13-8 lists the port error messages that can be displayed.

TABLE 13-8 Port Errors

Error Name	Numerical Value	Message
PRT_UNKNOWNPORT	0x300000	bad port number
PRT_ALREADYMAPPED	0x300001	port is already mapped unmap first
PRT_INVALIDNAME	0x300002	volume name is not correct

TABLE 13-8 Port Errors (Continued)

Error Name	Numerical Value	Message
PRT_VOLNOTFOUND	0x300003	volume name not found
PRT_INVALID	0x300004	port number is incorrect
PRT_LUNNOTMAPPED	0x300005	this lun is not mapped
PRT_ACCESSINVALID	0x300006	need to specify the access mode

Interconnect Card and Other FRU Errors

TABLE 13-9 lists the various FRU-related errors that you can encounter. These include various power and cooling unit fault conditions, missing disks, and interconnect card errors.

TABLE 13-9 Unit-Related Errors (Interconnect Card and Other FRUs)

Error Name	Numerical Value	Message
PS1_NOTEXIST	0x500001	Power Supply 1 Not Present
PS2_NOTEXIST	0x500002	Power Supply 2 Not Present
PS1_NOBATT	0x500003	Power Supply 1 Battery Not Present
PS2_NOBATT	0x500004	Power Supply 2 Battery Not Present
PS1_DISABLED	0x500005	Power Supply 1 Disabled
PS2_DISABLED	0x500006	Power Supply 2 Disabled
PS1_DISABLE	0x500007	Power Supply 1 Shutting off...
PS2_DISABLE	0x500008	Power Supply 2 Shutting off...
PS1_FAN1FAIL	0x500011	Power Supply 1 Fan 1 Failed
PS2_FAN1FAIL	0x500012	Power Supply 2 Fan 1 Failed
PS1_FAN2FAIL	0x500013	Power Supply 1 Fan 2 Failed
PS2_FAN2FAIL	0x500014	Power Supply 2 Fan 2 Failed
PS1_OVERTEMP	0x500015	Power Supply 1 Over Temperature
PS2_OVERTEMP	0x500016	Power Supply 2 Over Temperature
PS1_SWOFF	0x500017	Power Supply 1 Switch Off
PS2_SWOFF	0x500018	Power Supply 2 Switch Off
PS1_DCNOK	0x500019	Power Supply 1 DC Not OK
PS2_DCNOK	0x50001A	Power Supply 2 DC Not OK

TABLE 13-9 Unit-Related Errors (Interconnect Card and Other FRUs) *(Continued)*

Error Name	Numerical Value	Message
PS1_ONBATT	0x500021	Power Supply 1 On Battery
PS2_ONBATT	0x500022	Power Supply 2 On Battery
PS1_FANHIGH	0x500023	Power Supply 1 Fan High
PS2_FANHIGH	0x500024	Power Supply 2 Fan High
PS1_REFBATT	0x500025	Power Supply 1 Refresh Battery
PS2_REFBATT	0x500026	Power Supply 2 Refresh Battery
DK1_NOTEXIST	0x500031	Disk 1 Not Present
DK2_NOTEXIST	0x500032	Disk 2 Not Present
DK3_NOTEXIST	0x500033	Disk 3 Not Present
DK4_NOTEXIST	0x500034	Disk 4 Not Present
DK5_NOTEXIST	0x500035	Disk 5 Not Present
DK6_NOTEXIST	0x500036	Disk 6 Not Present
DK7_NOTEXIST	0x500037	Disk 7 Not Present
DK8_NOTEXIST	0x500038	Disk 8 Not Present
DK9_NOTEXIST	0x500039	Disk 9 Not Present
DK_NONE	0x50003A	No Disk Present
DK1_BYPASSED	0x500041	Disk 1 Bypassed
DK2_BYPASSED	0x500042	Disk 2 Bypassed
DK3_BYPASSED	0x500043	Disk 3 Bypassed
DK4_BYPASSED	0x500044	Disk 4 Bypassed
DK5_BYPASSED	0x500045	Disk 5 Bypassed
DK6_BYPASSED	0x500046	Disk 6 Bypassed
DK7_BYPASSED	0x500047	Disk 7 Bypassed
DK8_BYPASSED	0x500048	Disk 8 Bypassed
DK9_BYPASSED	0x500049	Disk 9 Bypassed
DK1_NOTREADY	0x500051	Disk 1 Not Ready
DK2_NOTREADY	0x500052	Disk 2 Not Ready
DK3_NOTREADY	0x500053	Disk 3 Not Ready
DK4_NOTREADY	0x500054	Disk 4 Not Ready
DK5_NOTREADY	0x500055	Disk 5 Not Ready

TABLE 13-9 Unit-Related Errors (Interconnect Card and Other FRUs) (Continued)

Error Name	Numerical Value	Message
DK6_NOTREADY	0x500056	Disk 6 Not Ready
DK7_NOTREADY	0x500057	Disk 7 Not Ready
DK8_NOTREADY	0x500058	Disk 8 Not Ready
DK9_NOTREADY	0x500059	Disk 9 Not Ready
CT_NOTEXIST	0x500061	Controller Not Present
CT_QLOGNRDY	0x500062	Qlogic Chip Not Ready
CT_SEL_ID	0x500063	Select ID Changed
LP_VSC_ERR	0x500064	VSC7120 Loop Failed
LC1_OFFLINE	0x500065	Loop Card 1 Offline
LC2_OFFLINE	0x500066	Loop Card 2 Offline
LP_CABLE1	0x500067	Cable 1 Not Present
LP_CABLE2	0x500068	Cable 2 Not Present
LC1_NSTART	0x500069	Loop Card 1 Failed to Start
LC2_NSTART	0x50006A	Loop Card 2 Failed to Start
CT_NOALTLP	0x50006B	No Alternate Loop
LP_SWITCH1	0x500071	Switch to Loop 1
LP_SWITCH2	0x500072	Switch to Loop 2
LP_MUX_ISO	0x500073	Loop Mux Changed to Isolated
LP_MUX_TOP	0x500074	Loop Mux Changed to Top
LP_MUX_MID	0x500075	Loop Mux Changed to Middle
LP_MUX_BOT	0x500076	Loop Mux Changed to Bottom

Other Errors

You will rarely see other types of errors, such as logical volume manager (LVM) errors (range 0x10001-0x1000A) and operating system errors (range 00000001-C000FFFF). The `tftp error` (numerical values 10060001-10060005) is an exception that you can see when you try to download a new binary. The `tftp errors` are usually generated due to one of the following reasons:

- The permissions for the file to be downloaded are too restrictive. In general, binaries should be world readable and executable.
- The checksum for the binary file to be downloaded is erroneous.

- The array units have not been recognized by the network. In this case, a system administrator should make sure that the IP addresses of the arrays are entered in the network database.

TABLE 13-10 lists pSOS errors.

TABLE 13-10 Embedded Operating System and Driver Errors

Error Type	Numerical Value
pSOS+	0000'0001 0000'0FFF
(reserved)	0000'1000 0000'1FFF
embedded file system	0000'2000 0000'2FFF
pREPC+	0000'3000 0000'3FFF
(reserved)	0000'4000 0000'4FFF
pNA+, pRPC+, pX11+	0000'5000 0000'5FFF
(reserved)	0000'6000 0000'FFFF
Networking libraries	0110'0000 01FF'FFFF
MMUlib	0120'0000 0120'00FF
(reserved)	0120'0100 1000'FFFF
serial driver	1001'0000 1001'FFFF
tick timer driver	1002'0000 1002'FFFF
(reserved)	1003'0000 1003'FFFF
RAM disk driver	1004'0000 1004'FFFF
(reserved)	1005'0000 1005'FFFF
TFTP driver	1006'0000 1006'FFFF
SLIP driver	1007'0000 1007'FFFF
(reserved)	1008'0000 1004'FFFF
SCSI driver	1050'0000 105F'FFFF
(reserved)	1060'0000 FFFF'FFFF

Examples

This section contains examples for different types of messages.

- Error Messages
- Warning Messages
- Notice

Error Messages

This section provides examples of the types of errors that may be encountered and the corresponding error message.

FRU Level Errors

- Example 1

When a FRU is removed, it must be replaced in less than 30 minutes or the unit will automatically shut down. In this example, the PCU in unit 1 is missing.

```
Jan 28 22:16:16 TMRT[1]: E: ulpcul: Missing; system shutting down in 25 minutes
Jan 28 22:21:16 TMRT[1]: E: ulpcul: Missing; system shutting down in 20 minutes
Jan 28 22:26:16 TMRT[1]: E: ulpcul: Missing; system shutting down in 15 minutes
Jan 28 22:31:16 TMRT[1]: E: ulpcul: Missing; system shutting down in 10 minutes
Jan 28 22:36:16 TMRT[1]: E: ulpcul: Missing; system shutting down in 5 minutes
Jan 28 22:41:16 TMRT[1]: E: ulpcul: Missing; system shutting down now
```

■ Example 2

A power and cooling unit (PCU) has an over-temperature condition.

```
E: u<n>pcu<n>: Over temperature
```

■ Example 3

The power and cooling unit (PCU) is flagged as not being present.

```
E: u<n>pcu<n>: Not present
```

Warning Messages

This section provides examples of the types of errors that may occur and the corresponding warning messages.

Unrecoverable Drive Error

Data drive `uld4` error occurred while using interconnect path 0.

```
Jan 25 00:09:20 ISR1[1]: W: uld4 SCSI Disk Error Occurred (path = 0x0)
Jan 25 00:09:20 ISR1[1]: W: Sense Key = 0x4, Asc = 0x15, Ascq = 0x1
Jan 25 00:09:20 ISR1[1]: W: Sense Data Description = Mechanical Positioning Error
```

This example shows the decode of the prior SCSI sense key information (from the data drive) regarding the drive error. In this case, the drive error was a “Mechanical Position Error.”

```
Jan 25 00:09:21 WXFT[1]: W: uld4 hard err in vol (n1) starting auto disable
```

This message indicates that drive `uld4` was disabled as a result of the previous error. The unit will continue to operate with a disabled drive, with possible performance degradation. A repair action must be planned. However, since the drive automatically disabled, the repair can be deferred.

Cache Parity Error

The Fibre Channel chip, model ISP2200 interconnect 1, on the controller card detected a parity error.

```
Dec 22 14:09:50 ROOT[1]: W: (ctr) SysFail Reset (7010) was initiated at 19981222
130800 ISP2100[1] detected parity error
```

In this example:

- W: (ctr) indicates that the controller card needs to be replaced.
- SysFail Reset (7010) indicates a reset log type.

This message points to a parity error occurring on the local bus of the controller. The error is examined to determine if it is a single-bit error. If it is a single-bit error, an attempt is made to rewrite the location. If the rewrite is successful, the error is viewed as temporary and a notice, containing the address, is written to the `syslog` with a unique identifier. If the rewrite is not successful, the error is considered permanent and a warning, containing the address, is sent upstream to the monitoring applications.

If another error occurs while operating with a permanent single-bit error, the controller will fail over to the other controller.

Software Assertion or Panic

Reset type.

```
Jan 30 03:01:48 ROOT[1]: W: ulctr Assertion Reset (3000) was initiated at
20000130 030050 g78/src/noblep/sdl/scsi/svd.c line 829, Assert(lid >=
SVD_START_LID) => 0 BOOT

Jan 30 03:01:48 ROOT[1]: N: CPU state...
Jan 30 03:01:48 ROOT[1]: N: R0 = 00000000 01554588 00250ea4 015545cc 00000000
ffffffff 01554634 00efe334
Jan 30 03:01:48 ROOT[1]: N: R8 = ffffffff 00002051 000d7438 00440000 00438f74
00252418 2fff0000 00000001
Jan 30 03:01:48 ROOT[1]: N: R16= 00000000 00000000 0023b828 00e40308 00e40308
00000008 00000009 00000000
Jan 30 03:01:48 ROOT[1]: N: R24= 00000000 2ffe0000 0023abe8 00060894 00000000
00efe324 00efe318 0023b940
```

This message points to a software assertion or panic on the controller. This error causes a controller panic. The controller is reset and a warning message is logged.

Unrecoverable Drive Error

```
Dec 22 16:57:55 ISR1[1] : W: (dr) SCSI Disk Error Occurred (path = 0x0,uld8)
```

- `W: (dr)` indicates that the disk drive requires repair.
- `path = 0x0` indicates that the error was detected while using the interconnect served by interconnect card 1 (`path 0x1` is interconnect served by interconnect card 2).
- `uld8` is the address of the failing drive.

```
Dec 22 16:57:55 ISR1[1] : W: (dr) Sense key = 0xb, Asc 0x47, Ascq = 0x0  
Dec 22 16:57:55 ISR1[1] : W: (dr) Sense Data Description = SCSI Parity Error
```

- `Sense key 0xb = ABORTED COMMAND`. Indicates that the TARGET aborted the command. The INITIATOR may be able to recover by retrying the command.
- The meaning of `Asc` and `Ascq` is decoded in the second line.
- `SCSI Parity Error` is the decode of the most important information contained in the `Sense key` message of the previous line.

These messages point to a parity error being detected and reported by the associated disk drive.

Recoverable Drive Error

```
Feb 04 22:39:11 ISR1[1] : W: uld1 SCSI Disk Error Occurred (path = 0x1)
```

- `W: uld1` indicates the failing drive.
- `path = 0x1` indicates that the error was detected while using the interconnect served by interconnect card 2 (`path 0x0` is interconnect served by interconnect card 1).

```
Feb 04 22:39:11 ISR1[1] : W: Sense key = 0x1, Asc 0x47, Ascq = 0x1  
Feb 04 22:39:11 ISR1[1] : W: Sense Data Description = Recovered Data With Retries
```

- `Sense key (N (hex)) = RECOVERED ERROR`. Indicates that the last command completed successfully after some recovery action was performed.
- The meaning of `Asc` and `Ascq` is decoded in the second line.

- Recovered Data With Retries is the decode of the most important information contained in the Sense key message of the previous line.

Notice

```
Feb 25 21:47:03 LPCT[1]: N: ulpcu2: Switch off
```

- N: is the Notice level message.
- ulpcu2 is unit 1, power and cooling unit 2.
- Switch off indicates that the power switch is off.

Glossary

A

- administrative domain** Partner groups (interconnected controller units) that share common administration through a master controller.
- alternate master controller unit** Also called “alternate master unit,” the secondary array unit in a partner group that provides failover capability from the master controller unit.
- Alternate Pathing (AP)** A mechanism that reroutes data to the other array controller in a partner group upon failure in the host data path. Alternate Pathing requires special software to perform this function.
- auto cache mode** The default cache mode for the Sun StorEdge T3 and T3+ array. In a fully redundant configuration, cache is set to write-behind mode. In a nonredundant configuration, cache is set to write-through mode. Read caching is always performed.
- auto disable** The Sun StorEdge T3 and T3+ array default that automatically disables a disk drive that has failed.

B

- buffering** Data that is being transferred between the host and the drives.

C

**command-line interface
(CLI)**

The interface between the Sun StorEdge T3 and T3+ array's pSOS operating system and the user in which the user types commands to administer the array.

controller unit

A Sun StorEdge T3 and T3+ array that includes a controller card. It can be used as a standalone unit or configured with other Sun StorEdge T3 and T3+ arrays.

D

**Dynamic Multi-Pathing
(DMP)**

A VERITAS Volume Manager feature that provides an Alternate Pathing mechanism for rerouting data in the event of a controller failover.

E

**enterprise
configuration**

One or more partner groups (pair of interconnected controller units) in a system configuration.

**erasable programmable
read-only memory
(EPROM)**

Memory stored on the controller card; useful for stable storage for long periods without electricity while still allowing reprogramming.

expansion unit

A Sun StorEdge T3 and T3+ array without a controller card. It must be connected to a controller unit to be operational.

F

**Fibre Channel
Arbitrated Loop
(FC-AL)**

A 100-Mbyte/s serial channel that enables connection of multiple devices (disk drives and controllers).

field-replaceable unit (FRU) A component that is easily removed and replaced by a field service engineer or a system administrator.

FLASH memory device (FMD) A device on the controller card that stores EPROM firmware.

G

Gigabit Interface Converter (GBIC) An adapter used on an SBus card to convert fiber-optic signal to copper.

gigabyte (GB or Gbyte) One gigabyte is equal to one billion bytes (1X10⁹).

graphical user interface (GUI) A software interface that enables configuration and administration of the Sun StorEdge T3 and T3+ array using a graphic application.

H

host bus adapter (HBA) An adapter that resides on the host.

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 system remains powered on and operational.

I

input/output operations per second (IOPS) A performance measurement of the transaction rate.

interconnect cable An FC-AL cable with a unique switched-loop architecture that is used to interconnect multiple Sun StorEdge T3 and T3+ arrays.

interconnect card An array component that contains the interface circuitry and two connectors for interconnecting multiple Sun StorEdge T3 and T3+ arrays.

L

- LC** An industry-standard name used to describe a connector standard. The Sun StorEdge T3+ array uses an LC-SFF connector for the host FC-AL connection.
- light-emitting diode (LED)** A device that converts electrical energy into light that is used to display activity.
- logical unit number (LUN)** One or more drives that can be grouped into a unit; also called a volume.

M

- master controller unit** Also called a “master unit,” the main controller unit in a partner-group configuration.
- media access control (MAC) address** A unique address that identifies a storage location or a device.
- media interface adapter (MIA)** An adapter that converts fiber-optic light signals to copper.
- megabyte (MB or Mbyte)** One megabyte is equal to one million bytes (1X106).
- megabytes per second (MB/s)** A performance measurement of the sustained data transfer rate.
- multi-initiator configuration** A supported array configuration that connects two hosts to one or more array administrative domains through hub or switch connections.

P

- parity** Additional information stored with data on a disk that enables the controller to rebuild data after a drive failure.
- partner group** A pair of interconnected controller units.

**power and cooling unit
(PCU)**

A FRU component in the Sun StorEdge T3 and T3+ array. It contains a power supply, cooling fans, and an integrated UPS battery. There are two power and cooling units in a Sun StorEdge T3 and T3+ array.

pSOS

An operating system built into the Sun StorEdge T3 and T3+ array RAID Controller firmware, which provides interfaces between the mounted RAID volumes and the database environment.

Q

quiesce To halt all drive activity.

R

read caching Data stored for future retrieval, to reduce disk I/O as much as possible.

**redundant array of
independent disks
(RAID)**

A configuration in which multiple drives are combined into a single virtual drive to improve performance and reliability.

**reliability, availability,
serviceability (RAS)**

A term to describe product features that include high availability, easily serviced components, and dependability.

**reverse address
resolution protocol
(RARP)**

A utility in the Solaris operating environment that enables automatic assignment of the array IP address from the host.

S

SC An industry standard name used to describe a connector standard.

**Simple Network
Management Protocol
(SNMP)**

A network management protocol designed to give a user the capability to remotely manage a computer network.

**small form factor
(SFF)**

An industry standard describing a type of connector. An LC-SFF connector is used for the host FC-AL connection to the Sun StorEdge T3+ array.

**synchronous dynamic
random access memory
(SDRAM)**

A form of dynamic random access memory (DRAM) that can run at higher clock speeds than conventional DRAM.

system area

Located on the disk drive label, the space that contains configuration data, boot firmware, and file-system information.

U

**uninterruptable power
source (UPS)**

A component within the power and cooling unit. It supplies power from a battery in the case of an AC power failure.

V

volume

Also called a logical unit or LUN, a volume is one or more drives that can be grouped into a unit for data storage.

W

**workgroup
configuration**

A standalone array connected to a host system.

**world wide name
(WWN)**

A number used to identify array volumes in both the array system and Solaris operating environment.

write caching Data used to build up stripes of data, eliminating the read-modify-write overhead. Write caching improves performance for applications that are writing to disk.

Index

A

- access permission, 18
- accessing Sun documentation online, xviii
- administration overview, 23
- advanced array troubleshooting, 113
- architecture, 12
- array
 - administration
 - overview, 23
 - to display command syntax, 24
 - to display commands, 24
 - characteristics, 2
 - configuration considerations, 73
 - firmware 2.1 new features, 16
 - monitoring
 - to check data parity, 48
 - to check the battery, 49
 - to check the drive status, 47
 - to check the hot-spare, 48
 - to determine failover, 46
 - to display FRU information, 50
 - to display FRU status, 51
 - troubleshooting, 113
- array configuration
 - for remote system logging
 - to edit the /etc/hosts file, 59
 - to edit the /etc/syslog.conf file, 58
 - to edit the management host /etc/syslog.conf file, 61
 - to transfer files back to the array, 60
 - to transfer files to the management host, 56
 - for SNMP notification

- to edit the array /etc/hosts file, 67
- to edit the array /etc/syslog.conf file, 65
- to transfer files back to the array, 68
- to transfer the array files to the management host, 63
- auto upgrade, 20
- auto-versioning, 21

B

- battery, checking, 49, 79

C

- cache
 - mode, setting, 30
 - read-ahead threshold, setting, 34
 - segment, 28
 - segment size
 - setting, 28
 - segment size, displaying, 29
- capacity per unit, 11
- CLI (command-line interface) command descriptions
 - code conventions, 135
 - complete list, 132
 - Fibre Channel topology, 185
 - FRU identifiers, 135
 - general T3+ command list, 136
 - LUN mapping, 165
 - LUN masking, 171

- Network Time Protocol (NTP), 189
 - volume slicing, 158
- command-line interface (CLI), xv
- commands
 - displaying all, 24
 - syntax, displaying, 24
- components, 4
 - controller card description, 5
 - disk drive description, 4
 - power and cooling unit description, 10
- configuration considerations, 73
- configuring the array
 - for remote system logging
 - to edit the /etc/hosts file, 59
 - to edit the /etc/syslog.conf file, 58
 - to edit the management host /etc/syslog.conf file, 61
 - to transfer files back to the array, 60
 - to transfer files to the management host, 56
 - for SNMP notification
 - to edit the array /etc/hosts file, 67
 - to edit the array /etc/syslog.conf file, 65
 - to transfer files back to the array, 68
 - to transfer the array files to the management host, 63
- connections
 - Ethernet, 7
 - FC-AL, 6
 - serial, 7
- controller card
 - description, 5
 - replacement, 123
 - Sun StorEdge T3+ array controller, 7
- controller card, replacement, 123

D

- data, 40
 - block size definition, 28
 - cache SDRAM, 7
 - parity, checking, 48
- default slice, 16
- disk drive
 - description, 4
 - status messages, 47
 - status, checking, 47

- system
 - area, 12

E

- Ethernet connection, 7

F

- fabric point-to-point, 20
- fail over, determining, 46
- features
 - administration, 12
 - capacity per unit, 11
 - performance, 11
 - RAID capabilities, 12
 - reliability and redundancy, 12
- Fibre Channel-Arbitrated Loop (FC-AL), 6
- format utility, 42
- FRU (field-replaceable unit)
 - identifiers, 25
 - identifiers with commands, 25
 - information, displaying, 50
 - status, displaying, 51

G

- global parameters, 27

H

- host multipathing, 95
- hot spare
 - checking, 48
 - configuration limitations, 38

I

- input/output operations per second (IOPS), 11
- installation
 - verifying the firmware level, 119
- interconnect cards
 - description, 7 to 9

IP address, setting, 35

L

list of CLI commands, 132

Logical, 39

LUN

mapping

to add an entry to the LUN map, 90

to display the LUN map, 89

to remove an entry from the LUN map, 91

masking, 18

to change the default LUN permissions, 96

to define a host WWN group, 97

to find the default LUN permissions and system-assigned WWNs, 93

to remove a WWN from a specific group, 101

to remove the registered status of a WWN, 100

to set a specific LUN permission, 94

to set LUN permission with host multipathing, 95

to set the permissions of a LUN for all members of a group, 98

to view all registered WWNs, 94

permissions, 18

reconstruction rate, setting, 32

M

MAC address location, 127

masking a LUN, 92

messages

error example, 230

notice example, 231

warning example, 231

mirrored cache, enabling, 31

monitoring the array

to check data parity, 48

to check the battery, 49

to check the drive status, 47

to check the hot-spare, 48

to determine failover, 46

to display FRU information, 50

to display FRU status, 51

multipathing software, 95

multiple volume slicing, 18

N

new features, 2

auto upgrade, 20

auto-versioning, 21

fabric point-to-point, 20

list, 16

network time protocol, 21

persistent group reservation, 20

split loop mode, 20

volume slicing, mapping, and masking, 16

notice message

example, 234

NTP (network time protocol), 21

O

online Sun documentation, xviii

ordering Sun documentation, xix

overview of product, 1

P

partner group multipathing, enabling, 30

performance, 11

permission, 18

PGR (persistent group reservation), 20

point-to-point, 20

power and cooling unit

description, 10

replacement, 123

product

description, 2

overview, 1

R

RAID capabilities, 12

RAID level

configuration limitations, 38

considerations, 74

- RAS (reliability, availability, and serviceability)
 - features, 2
- reconfiguring a volume
 - creating, 39
 - deleting, 38
 - labeling, 41
- reconstruction rates, 32
- related documentation, xvii
- reliability and redundancy, 12
- reliability, availability, and serviceability (RAS)
 - features, 2
- reset
 - firmware version 2.0.1, 77
 - firmware version 2.1, 77
- resetting global parameters
 - to enable mirrored cache, 31
 - to enable partner group multipathing, 30
 - to perform volume verification, 33
 - to set cache read-ahead threshold, 34
 - to set the cache block size, 28
 - to set the cache mode, 30
 - to set the IP address, 35
 - to set the LUN reconstruction rate, 32

S

- scalability of array, 11
- SDRAM data cache, 7
- serial number location, 127
- shell prompts, xviii
- slice, default, 16
- slicing
 - a volume, 75
 - limitations, 73
- split loop mode, 20
- stripe unit size
 - See data block size
- Sun documentation
 - online accessing, xviii
 - ordering, xix
- Sun StorEdge T3+ array
 - architecture, 12
 - characteristics, 2
 - components, 4
 - controller card and ports, 7
 - description, 2

- features, 11
 - new features, 2
 - scalability, 11
- sys command
 - blocksize subcommand, 29
 - determining fail over, 46
 - enabling multipathing, 30
 - mirrored cache, 31
 - read-ahead threshold, 35
 - reconstruction rates, 33
- syslog error messages
 - command-line error message descriptions, 220
 - interconnect card and other FRU errors, 226
 - other errors, 228
 - port errors, 225
 - RAID and other errors, 221
 - error message descriptions, 203
 - error message syntax, 202
 - examples, 230
 - message types, 202
 - reset log types, 219
 - warning message descriptions, 208
- system
 - area, 12
 - logging, configuring, 51, 63

T

- troubleshooting
 - replacing FRUs, 119
 - to identify a data channel failure, 130
 - to identify miscabled partner groups, 126
 - to install a replacement FRU, 121
 - to rebuild a replaced drive FRU, 121
 - to verify the FRU firmware level, 119
 - to verify the FRU replacement, 124
- typographical conventions, xvi

U

- UNIX commands, using, xviii

V

- verification frequency, setting, 33

- vol command
 - adding a volume, 40
 - checking data parity, 48
 - initializing volumes, 40
 - list subcommand, 28
 - mounting volume, 40
 - rebuilding a replaced fru, 121
 - remove subcommand, 28
 - unmount subcommand, 28
 - verification frequency, 34
 - verify subcommand, 48
- volume
 - configuration limitations, 72
 - initializing, 40
 - mounting, 40
 - performing verification, 33
 - reconfiguration limitations, 38
 - reconfiguring
 - creating, 39
 - deleting, 38
 - labeling, 41
- volume slicing, 18
 - limitations, 73
 - to create a volume slice, 88
 - to enable volume slicing, mapping, and masking, 76
 - to reconfigure a volume by adding slices after volume slicing has been enabled, 79
 - to remove a volume slice, 88

W

- WWN (world-wide number), 19

