



Sun StorEdge™ T3+ Array Installation and Configuration Manual

Version 2.1 Controller Firmware

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

Part No. 816-4769-10
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, SunSolve, JumpStart, StorTools, Solstice Backup, SunService, OpenBoot, Sun Enterprise, Sun Fire, Sun VTS, Solstice DiskSuite, Sun StorEdge, 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, SunSolve, JumpStart, StorTools, Solstice Backup, SunService, OpenBoot, Sun Enterprise, Sun Fire, Sun VTS, Solstice DiskSuite, Sun StorEdge, 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 xvii

- 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 9
 - Array Features 9
 - Firmware Version 2.1 New Features List 11
 - Array Architecture 11
- 2. Supported Configurations 13**
 - Workgroup 14
 - Enterprise 16
 - Configuration Guidelines and Restrictions 17
 - Workgroup Configuration Restrictions 17
 - Enterprise Configuration Restrictions 17
 - Configuration Recommendations 18

Sun Cluster Support 19

3. Supported Software 21

Supported Operational Environment 22

Supported Multiple Operational Environments 22

Supported Multi-Pathing Software 22

Supported Clustering Software 22

Supported Host Side Drivers 23

Supported SAN Switch Support 23

Supported Diagnostics and Monitoring Software 23

 Administrative Tools 23

 End-User Tools 24

Supported Interfaces 24

Supported Import Interfaces 24

Supported Bootability 25

Supported Multi-Platform Software 25

Supported Multi-Host Environments 25

Summary of Supported Software 26

4. Configuration Scenarios 27

Direct Host Connection 28

 Single Host With One Controller Unit 28

 Single Host With Two Controller Units Configured as a Partner Group 29

 Host Multi-Pathing Management Software 30

 Single Host With Four Controller Units Configured as Two Partner Groups 31

 Single Host With Eight Controller Units Configured as Four Partner Groups 33

Hub Host Connection 35

| | |
|--|-----------|
| Single Host With Two Hubs and Four Controller Units Configured as Two Partner Groups | 35 |
| Single Host With Two Hubs and Eight Controller Units Configured as Four Partner Groups | 38 |
| Dual Hosts With Two Hubs and Four Controller Units | 40 |
| Dual Hosts With Two Hubs and Eight Controller Units | 42 |
| Dual Hosts With Two Hubs and Four Controller Units Configured as Two Partner Groups | 44 |
| Dual Hosts With Two Hubs and Eight Controller Units Configured as Four Partner Groups | 46 |
| Switch Host Connection | 48 |
| Dual Hosts With Two Switches and Two Controller Units | 48 |
| Dual Hosts With Two Switches and Eight Controller Units | 50 |
| 5. Preparing to Install | 53 |
| Preparing for the Installation | 54 |
| ▼ To Inspect the Array | 55 |
| ▼ To Edit the Host Files | 57 |
| ▼ To Install Required Patches | 59 |
| Installation Checklist | 61 |
| 6. Array Installation | 63 |
| ▼ To Install the Array | 64 |
| ▼ To Connect the Cables | 65 |
| ▼ To Power On and Verify the Hardware Configuration | 70 |
| ▼ To Establish a Network Connection | 72 |
| ▼ To Verify the Firmware Level and Configuration | 75 |
| Defining and Mounting Volumes | 81 |
| ▼ To Configure the Workgroup | 83 |
| ▼ To Configure the Enterprise | 85 |
| Changing the Default Configuration | 88 |

| | |
|--|------------|
| Connecting the Host System | 89 |
| Verifying the Data Host Connection | 89 |
| Establishing Logical Volumes on the Host | 90 |
| Creating Soft Partitions on the Data Host | 90 |
| ▼ To Establish Alternate Pathing on the Host | 91 |
| Powering Off and On | 94 |
| ▼ To Power Off | 94 |
| ▼ To Power On | 95 |
| 7. Controller Firmware Version 2.1 Upgrade | 97 |
| Upgrading Overview | 98 |
| Upgrading Requirements | 98 |
| Upgrading Limitations | 99 |
| Connecting to the Array | 100 |
| ▼ To Establish an ftp Session | 100 |
| ▼ To Establish a Telnet Session | 101 |
| ▼ To Determine the Firmware Version | 103 |
| Upgrading the System Firmware | 104 |
| ▼ To Install the System Firmware 2.1 Patch | 104 |
| ▼ To Install the Interconnect Card Firmware | 107 |
| ▼ To Upgrade the Controller Firmware | 109 |
| ▼ To Verify the Upgrade | 111 |
| 8. Post-Installation Global Parameters Configurations | 115 |
| Configuring Cache | 116 |
| Performance and Redundancy | 116 |
| Data Block Size Configuration | 117 |
| Data Block Size Selection | 117 |
| Mirrored Cache Enabling | 118 |

| | |
|---|------------|
| Cache Allocation Configuration | 118 |
| Configuring Logical Volumes | 119 |
| Guidelines for Configuring Logical Volumes | 119 |
| Determining the Number of Logical Volumes | 120 |
| Determining the RAID Level | 120 |
| Determining a Hot-Spare | 121 |
| Creating and Labeling a Logical Volume | 122 |
| Setting the LUN Reconstruction Rate | 122 |
| Using RAID Levels to Configure Redundancy | 123 |
| RAID 0 | 124 |
| RAID 1 | 124 |
| RAID 5 | 125 |
| Configuring RAID Levels | 125 |
| 9. Post-Installation Array Settings Configuration | 127 |
| Reconfiguring the Array Settings | 128 |
| Unit Volume Configurations | 128 |
| RAID Levels | 129 |
| Cache Modes | 129 |
| Disabling and Reconstructing a Drive | 131 |
| Monitoring the Array | 132 |
| SNMP Notification | 132 |
| syslog Error Reports | 133 |
| 10. Post-Installation Partner Groups Configuration | 135 |
| Understanding Partner Groups | 136 |
| How Partner Groups Work | 137 |
| Creating Partner Groups | 138 |
| 11. Post-Installation Host Connections | 139 |

| | |
|--|------------|
| Sun Enterprise SBus+ and Graphics+ I/O Boards | 140 |
| System Requirements | 140 |
| Sun StorEdge PCI FC-100 Host Bus Adapter | 141 |
| System Requirements | 141 |
| Sun StorEdge SBus FC-100 Host Bus Adapter | 142 |
| System Requirements | 142 |
| Sun StorEdge PCI Single Fibre Channel Network Adapter | 143 |
| System Requirements | 143 |
| Sun StorEdge PCI Dual Fibre Channel Network Adapter | 144 |
| System Requirements | 144 |
| Compact PCISun StorEdge CompactPCI Dual Fibre Channel Network Adapter | 145 |
| System Requirements | 145 |
| 12. Post-Installation Array Cabling | 147 |
| Overview of Array Cabling | 148 |
| Data Path | 148 |
| Administration Path | 148 |
| Connecting Arrays | 148 |
| Workgroup Configurations | 150 |
| Enterprise Configurations | 151 |
| 13. Troubleshooting the Array | 153 |
| ofdg Notice | 154 |
| Analyzing the Problem | 155 |
| Host-Generated Message | 155 |
| Telnet Session | 155 |
| Array LEDs | 156 |
| Disk Drive LEDs | 156 |

| | | |
|------------|--|------------|
| | Power and Cooling Unit LEDs | 157 |
| | Interconnect Card LEDs | 159 |
| | Controller Card LEDs | 160 |
| | Channel Connection Failures | 161 |
| | FRU Failures | 161 |
| 14. | Servicing the Array | 163 |
| | ofdg Notice | 164 |
| | Preparing for Service | 165 |
| | Removing and Replacing Components | 165 |
| | Disk Drives | 166 |
| | Disk Drive LEDs | 166 |
| | ▼ To Remove and Replace a Disk Drive | 167 |
| | Power and Cooling Units | 170 |
| | Power and Cooling Unit LEDs | 171 |
| | ▼ To Remove and Replace a Power and Cooling Unit | 171 |
| | UPS Battery | 173 |
| | Battery Maintenance | 173 |
| | Removing and Replacing the UPS Battery | 174 |
| | Interconnect Cards | 174 |
| | Interconnect Card LEDs | 174 |
| | ▼ To Removing and Replacing an Interconnect Card | 175 |
| | Controller Card | 177 |
| | Controller Card LEDs | 177 |
| | ▼ To Remove and Replace a Controller Card | 177 |
| | Chassis | 179 |
| A. | Specifications | 181 |
| | Power Specifications | 182 |

Environmental Specifications 183

Mechanical Specifications 184

Cable Specifications 185

B. Illustrated Parts 187

Glossary 193

Figures

| | | |
|-------------|--|----|
| FIGURE 1-1 | Sun StorEdge T3+ Array | 3 |
| FIGURE 1-2 | Disk Drives (Front View) | 4 |
| FIGURE 1-3 | Controller Card (Rear View) | 5 |
| FIGURE 1-4 | Sun StorEdge T3+ Array Controller Card Front Panel | 5 |
| FIGURE 1-5 | Sun StorEdge T3+ array Controller Card and Ports | 6 |
| FIGURE 1-6 | Interconnect Card and Ports | 7 |
| FIGURE 1-7 | Interconnect Cards (Rear View) | 8 |
| FIGURE 1-8 | Power and Cooling Units (Rear View) | 9 |
| FIGURE 2-1 | Workgroup Configuration | 14 |
| FIGURE 2-2 | Enterprise Configuration | 15 |
| FIGURE 4-1 | Single Host Connected to One Controller Unit | 28 |
| FIGURE 4-2 | Single Host With Two Controller Units Configured as a Partner Group | 29 |
| FIGURE 4-3 | Failover Configuration | 30 |
| FIGURE 4-4 | Single Host With Four Controller Units Configured as Two Partner Groups | 32 |
| FIGURE 4-5 | Single Host With Eight Controller Units Configured as Four Partner Groups | 34 |
| FIGURE 4-6 | Single Host With Two Hubs and Four Controller Units Configured as Two Partner Groups | 37 |
| FIGURE 4-7 | Single Host With Two Hubs Configured and Eight Controller Units as Four Partner Groups | 39 |
| FIGURE 4-8 | Dual Hosts With Two Hubs and Four Controller Units | 41 |
| FIGURE 4-9 | Dual Hosts With Two Hubs and Eight Controller Units | 43 |
| FIGURE 4-10 | Dual Hosts With Two Hubs and Four Controller Units Configured as Two Partner Groups | 45 |

| | | |
|-------------|---|-----|
| FIGURE 4-11 | Dual Hosts With Two Hubs and Eight Controller Units Configured as Four Partner Groups | 47 |
| FIGURE 4-12 | Dual Hosts With Two Switches and Two Controller Units | 49 |
| FIGURE 4-13 | Dual Hosts With Two Switches and Eight Controller Units | 51 |
| FIGURE 5-1 | Sun StorEdge T3+ Array Cables | 55 |
| FIGURE 5-2 | Removing the Front Panel | 56 |
| FIGURE 5-3 | Serial Number and MAC Address on Pull-out Tab | 56 |
| FIGURE 6-1 | Sun StorEdge T3+ Array Cables | 66 |
| FIGURE 6-2 | Connecting the Fiber-Optic Cable to the FC-AL Connector on a Sun StorEdge T3+ Array | 67 |
| FIGURE 6-3 | Connecting the Ethernet Cable | 67 |
| FIGURE 6-4 | Connecting the Power Cords | 68 |
| FIGURE 6-5 | Connecting the Interconnect Cables | 69 |
| FIGURE 6-6 | Fully Cabled Partner Group (Enterprise Configuration) | 70 |
| FIGURE 6-7 | Power and Cooling Unit LEDs and Power Button Location | 71 |
| FIGURE 6-8 | Disk Drive Physical Numbering | 81 |
| FIGURE 6-9 | Power Switch Locations | 95 |
| FIGURE 7-1 | Level 3 Firmware Version Check Output for Master Unit | 103 |
| FIGURE 7-2 | Level 3 Firmware Version Check Output for Master Unit | 111 |
| FIGURE 10-1 | Sun StorEdge T3+ Array Partner Group | 136 |
| FIGURE 11-1 | Sun Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O Board | 140 |
| FIGURE 11-2 | Sun StorEdge PCI FC-100 Host Bus Adapter | 141 |
| FIGURE 11-3 | Sun StorEdge SBus FC-100 Host Bus Adapter | 142 |
| FIGURE 11-4 | Sun StorEdge PCI Single Fibre Channel Network Adapter | 143 |
| FIGURE 11-5 | Sun StorEdge PCI Dual Fibre Channel Network Adapter | 144 |
| FIGURE 11-6 | Sun StorEdge CompactPCI Dual Fibre Channel Network Adapter | 145 |
| FIGURE 12-1 | Sun StorEdge T3+ Array Controller Card and Interconnect Cards | 149 |
| FIGURE 12-2 | Array Workgroup Configuration | 150 |
| FIGURE 12-3 | Enterprise Configuration | 151 |
| FIGURE 13-1 | Disk Drive LEDs (Viewed Through Front Cover) | 156 |
| FIGURE 13-2 | Power and Cooling Unit LEDs | 157 |

| | | |
|-------------|--|-----|
| FIGURE 13-3 | Interconnect Card LEDs | 159 |
| FIGURE 13-4 | Sun StorEdge T3+ Array Controller Card LEDs | 160 |
| FIGURE 14-1 | Removing the Front Panel | 167 |
| FIGURE 14-2 | Disk Drive Numbering | 167 |
| FIGURE 14-3 | Releasing the Latch Handle | 168 |
| FIGURE 14-4 | Removing a Disk Drive | 169 |
| FIGURE 14-5 | Power Cords Connected to the Power and Cooling Units | 170 |
| FIGURE 14-6 | Removing the Power and Cooling Unit | 172 |
| FIGURE 14-7 | Removing the Interconnect Card | 176 |
| FIGURE 14-8 | Removing the Controller Card | 178 |
| FIGURE B-1 | Disk Drive in Canister | 188 |
| FIGURE B-2 | Sun StorEdge T3+ Array Controller Card | 189 |
| FIGURE B-3 | Interconnect Card | 189 |
| FIGURE B-4 | Power and Cooling Unit | 189 |
| FIGURE B-5 | Cables, Adapters, and Cords | 190 |
| FIGURE B-6 | Door Assembly (Front Panel) | 191 |

Tables

| | | |
|------------|---|-----|
| TABLE 1-1 | Sun StorEdge T3+ Array Key Features—73-Gbyte Drive | 10 |
| TABLE 5-1 | Required Patches | 60 |
| TABLE 6-1 | Drive Status Messages | 83 |
| TABLE 7-1 | Explanation of <code>ver</code> Command Results | 103 |
| TABLE 7-2 | <code>t3.sh</code> Utility Installation Prompts | 106 |
| TABLE 7-3 | Explanation of <code>ver</code> Command Results | 111 |
| TABLE 8-1 | RAID 0 Data Blocks | 124 |
| TABLE 8-2 | RAID 1 (1+0) Data Blocks | 124 |
| TABLE 8-3 | RAID 5 Data Blocks | 125 |
| TABLE 9-1 | Examples of Volume Configurations | 129 |
| TABLE 9-2 | Cache Transition States | 130 |
| TABLE 9-3 | <code>syslog</code> Error Reporting Levels | 133 |
| TABLE 13-1 | Disk Drive LED Descriptions | 157 |
| TABLE 13-2 | Power and Cooling Unit LED Descriptions | 158 |
| TABLE 13-3 | Interconnect Card LED Descriptions | 159 |
| TABLE 13-4 | Sun StorEdge T3+ Array Controller Card LED Descriptions | 160 |
| TABLE A-1 | Power Specifications | 182 |
| TABLE A-2 | Environmental Specifications | 183 |
| TABLE A-3 | Physical Specifications | 184 |
| TABLE A-4 | Cable Specifications | 185 |

Preface

This manual contains information on how to install, operate, and service the Sun StorEdge™ T3+ array. Most of the procedures in this manual can be performed by an experienced system administrator.

Topics covered here are:

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

Before You Read This Book

Read the safety information in the *Sun StorEdge T3+ Array Regulatory and Safety Compliance Manual* that is included with the array before attempting the installation. Make sure you have prepared for the installation by reviewing the *Sun StorEdge T3+ Array Site Preparation Guide*. These guides are available on the Sun™ documentation web site or from your Sun representative.

Work with your Sun representative to determine if you require any external hardware or software products for using this device. Being prepared with this knowledge and the appropriate tools will simplify the installation.

Using UNIX Commands

This document contains some information on basic UNIX® commands and procedures such as booting the devices. For further information, see one or more of the following:

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

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> . |

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 | # |
| Sun StorEdge T3 and T3+ array | :/ : |

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 |

Accessing Sun Documentation Online

You can find the Sun StorEdge T3+ array documentation and other select product documentation for Network Storage Solutions at:

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

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 (8xx-xxxx-xx) 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 9
- “Firmware Version 2.1 New Features List” on page 11
- “Array Architecture” on page 11

Product Description

The Sun StorEdge T3+ array 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 shown in FIGURE 1-1 can be placed on a table top or rackmounted in a server cabinet or expansion cabinet.

The array is sometimes called a *controller unit*, which refers to the internal RAID controller on the controller card. When connected to a controller unit, the expansion unit enables you to increase your storage capacity without the cost of an additional controller. An expansion unit must be connected to a controller unit to operate because it does not have its own controller.

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

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

- Fibre Channel disk storage solution
- Up to nine 73-GByte drives in each array
- RAID levels 0, 1, and 5 support
- Up to 1 Gbyte of data cache memory
- 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 Chapter 14.

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 Chapter 13 for more information on LED locations and descriptions. See Chapter 14 for information on how to service the disk drives.




FIGURE 1-2 Disk Drives (Front View)

Note – All Sun StorEdge T3+ array disk drives in a workgroup or in an enterprise configuration must have the same storage capacity and drive speed.

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. See FIGURE 1-3. The general features of each array controller card are described in this section.




FIGURE 1-3 Controller Card (Rear View)

The location of the controller card for both models is shown in FIGURE 1-4. Chapter 14 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 that uses 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 via 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. See FIGURE 1-7. For information on how to service the interconnect card, see Chapter 14.




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. Each array contains two power and cooling units for redundancy. See Chapter 14 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 | 657 GB (9 x 73-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 or 2 host Fibre Channel interfaces 90- to 190-MB/sec bandwidth 7,200 to 14,360 IOPS (100% cache hit reads) |
| 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 access (CLI interface) Simple Network Management Protocol (SNMP) Version 1 |

Firmware Version 2.1 New Features List

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

- Volume slicing
- LUN masking
- LUN mapping
- Auto upgrade
- Split loop mode
- Host side initiators
- Fabric point-to-point support
- Persistent group reservation (PGR)
- Network time protocol (NTP)

See the *Sun StorEdge T3+ Array Administrator's Manual* for more information on these features.

Array Architecture

The Sun StorEdge T3+ array is designed to be a modular, scalable, reliable, serviceable, high-performance storage building block with a flexible configuration. 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. *Controller units* have a cached hardware RAID controller for high performance. Controller units can be paired in an *enterprise (partner group)* configuration, providing controller and data path redundancy and mirrored caches for high availability. 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, of which approximately 150 Mbytes is reserved for the *system area*, 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, and enables 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
- Networked-based and serial-based (RS-232) administrative path
- Two separate serial buses to support control and sense operations.

Supported Configurations

This chapter describes the two supported array configurations: workgroup and enterprise. It also describes basic rules and recommendations for configuring the array.

Topics in this chapter include:

- “Workgroup” on page 14
- “Enterprise” on page 16
- “Configuration Guidelines and Restrictions” on page 17
- “Configuration Recommendations” on page 18
- “Sun Cluster Support” on page 19

Workgroup

This standalone array is a high-performance, high-RAS configuration with a single hardware RAID-cached controller. See FIGURE 2-1. The unit is fully populated with redundant *hot-swap* components and nine disk drives.




FIGURE 2-1 Workgroup Configuration




FIGURE 2-2 Enterprise Configuration

Also called a *partner group*, this is a configuration of two controller units paired with interconnect cables for back-end data and administrative connections. The *enterprise* configuration provides all the RAS of single controller units, plus redundant hardware RAID controllers with mirrored caches, and redundant host channels for continuous data availability for host applications.

Note – The Sun StorEdge T3+ array configuration supports direct FC-AL connections.

Enterprise

An enterprise configuration has a *master controller unit* and an *alternate master controller unit*. In all default enterprise configurations, the master controller unit is the array positioned at the bottom of an array stack in either a rack-mounted or tabletop installation. The alternate master controller unit is positioned on top of the master controller unit. The positioning of the master and alternate master controller units is important for cabling the units together correctly, understanding IP address assignments, interpreting array command-line screen output, and determining controller failover and failback conditions.

In this manual, the terms *master controller unit* and *master unit* are used interchangeably, as are the terms *alternate master controller unit* and *alternate master unit*.

Configuration Guidelines and Restrictions

Workgroup Configuration Restrictions

- The media access control (MAC) address is required to assign an IP address to the controller unit. The MAC address uniquely identifies each node of a network. The MAC address is available on the pull-out tab on the front left side of the array.
- A host-based mirroring solution is necessary to protect data in cache.
- Sun StorEdge T3+ array workgroup configurations are supported in Sun Cluster 3.0 environments.

Enterprise Configuration Restrictions

- Partner groups can be connected to more than one host only if all of the following conditions exist:
 - The partner group can be connected to the hosts through a hub or a switch.
 - The configuration is using Sun StorEdge Traffic Manager Software (STMS) or VERITAS Dynamic Multipathing (DMP) for multipathing support.
 - Cluster configurations are supported using Sun Cluster 3.0 or third-party software.
- You cannot use a daisy-chain configuration to link more than two controller units together.
- You can only connect arrays of the same model type in a partner group.
- In a cluster configuration, partner groups are supported only by Sun Cluster 3.0 software. They are not supported with Sun Cluster 2.2 software.



Caution – In an enterprise configuration, make sure you use the MAC address of the *master* controller unit.

Configuration Recommendations

- Use enterprise configurations for controller redundancy.
- Use host-based software such as VERITAS Volume Manager (VxVM), VERITAS Dynamic Multipathing (DMP), or Sun StorEdge Traffic Manager for multipathing support.
- Connect redundant paths to separate host adapters, I/O cards, and system buses.
- Configure active paths over separate system buses to maximize bandwidth.



Caution – For optimum performance, the array and its global parameters must be tailored to match the I/O workload. Within a partner group, both units will share the same *volume* configuration, block size, and cache mode. That is, all cache parameter settings are common to both units within a partner group.

- The Sun StorEdge SBus FC-100 (SOC+) HBA and the on-board SOC+ interface in Sun Fire™ systems are supported.
- On Sun Enterprise™ 6x00/5x00/4x00/3x00 systems, a maximum of 64 arrays are supported per cluster.
- On Sun Enterprise 10000 systems, a maximum of 256 arrays are supported per cluster.
- , To ensure full redundancy in a Sun StorEdge T3+ array workgroup configuration, host-based mirroring software such as Solstice Disk Suite (SDS) 4.2 and later or VERITAS Volume Manager *must* be used.
- Solaris 2.6 and Solaris 8 are the only supported operating systems.

Note – Refer to the latest Sun Cluster documentation for more information on array configurations and restrictions supported by Sun Cluster.

Sun Cluster Support

The Sun StorEdge T3+ array is supported in Sun Cluster configurations with the following restrictions:

- Array controller firmware version 2.0 or later is required on each Sun StorEdge T3+ array.
- Sun Cluster 3.0 environments support the Sun StorEdge T3+ array.
- Enterprise configurations are supported only in Sun Cluster 3.0 environments.
- Partner groups in a Sun Cluster environment must use Sun StorEdge Traffic Manager software for multipathing support.

Supported Software

Note – For the latest information on supported software, refer to the storage solutions web site at <http://www.sun.com/storage> and look for details on the Sun StorEdge T3+ array product family.

Topics in this chapter include:

- “Supported Operational Environment” on page 22
- “Supported Multiple Operational Environments” on page 22
- “Supported Clustering Software” on page 22
- “Supported Host Side Drivers” on page 23
- “Supported SAN Switch Support” on page 23
- “Supported Diagnostics and Monitoring Software” on page 23
- “Supported Interfaces” on page 24
- “Supported Import Interfaces” on page 24
- “Supported Bootability” on page 25
- “Supported Multi-Platform Software” on page 25
- “Summary of Supported Software” on page 26

Supported Operational Environment

- Solaris 2.6, and 8 operating environment software
- Microsoft NT operating environment software
- HP-UX
- AIX

Supported Multiple Operational Environments

With the volume slicing feature enabled, the slices can be masked to support multiple operating systems simultaneously.

Note – In the absence of a clustering software, managing each slice by a single host will be assumed.

Supported Multi-Pathing Software

- Veritas Volume Manager 3.2, 3.1.1
- Solaris Logical Volume Manger 4.2.1
- Solaris DiskSuite™
- Sun StorEdge Traffic Manager Software (STMS)

Supported Clustering Software

- Sun Cluster 2.2 and 3.0
- VERITAS Cluster Server

Supported Host Side Drivers

- SOC
- ifp
- Sun StorEdge Network Foundation Software
- Sun StorEdge Traffic Manager Software (STMS)

Supported SAN Switch Support

- Sun StorEdge 16-port FC-AL switch
- Sun StorEdge FC-AL switch

Supported Diagnostics and Monitoring Software

Administrative Tools

- Device Manager
- StorADE 2.0
- Network Storage Agent (Solaris 2.6 only)
- StorTools™ (Solaris 2.6 only)
- fru stat
- Ofdg
- Serial console
- SunVTS™

End-User Tools

- StorTools 3.1, 4.0, 4.0.1, 4.1
- StorADE 1.1
- Network Storage Agent 1.1
- fru stat
- Command-Line Interface (CLI)

Supported Interfaces

The protocols by which Sun StorEdge T3 2.1 firmware incorporates with others are:

SNMP-read-only access

HTTP-Token-based interface for use by Gotham

CLI-Command-line interface

Management applications use out-of-band interfaces with SNMP, TCP/IP, SMTP, HTTP protocols

Supported Import Interfaces

Fibre Channel protocol (specified in the ANSI FCP and FCP-2 specifications).

Device discovery and basic I/O operations are compliant with the Fibre Channel standards. T3+ array is a block storage

Device and uses FCP and FCP-2 protocol for interpretability compliance.

ANSI SCSI III Block Device Command Set Specification.

Supported Bootability

To use the Sun StorEdge T3+ array as a bootable device, contact your SunService™ provider for information.

Supported Multi-Platform Software

- T3 Multi-Platform 2.0

Supported Multi-Host Environments

- CSI 2 Reserve/Release
- SCSI III persistent group reservation (PGR)
- LUN masking

Summary of Supported Software

The following software is supported on Sun StorEdge T3+ array:

- Solaris 2.6, 8, and 9 operating environments
- VERITAS Volume Manager 3.1.1 and later with Dynamic Multipathing (DMP)
- Sun Enterprise Server Alternate Pathing (AP) 2.3.1
- StorTools 3.3 Diagnostics
- Sun Cluster 2.2 and 3.0 software (see “Sun Cluster Support” on page 19)
- Sun StorEdge Data Management Center 3.0
- Sun StorEdge Instant Image 2.0
- Sun StorEdge Network Data Replicator (SNDR) 2.0
- Solstice Backup™ 5.5.1
- Solstice DiskSuite 4.2 and 4.2.1

Configuration Scenarios

This chapter includes sample reference configurations for Sun StorEdge T3+ arrays. Although many configurations are supported, these reference configurations provide the best solution for many installations.

Topics in this chapter include:

- “Direct Host Connection” on page 28
- “Hub Host Connection” on page 35
- “Switch Host Connection” on page 48

Direct Host Connection

This section contains examples of the following configurations:

- “Single Host With One Controller Unit” on page 28
- “Single Host With Two Controller Units Configured as a Partner Group” on page 29
- “Single Host With Four Controller Units Configured as Two Partner Groups” on page 31
- “Single Host With Eight Controller Units Configured as Four Partner Groups” on page 33

Single Host With One Controller Unit

FIGURE 4-1 shows one application host connected through an FC-AL cable to one array controller unit. The Ethernet cable connects the controller to a management host by way of a secure private network, and requires an IP address.

Note – This configuration is not recommended for RAS functionality because the controller is a single point of failure. In this type of configuration, use a host-based mirroring solution to protect data in cache.

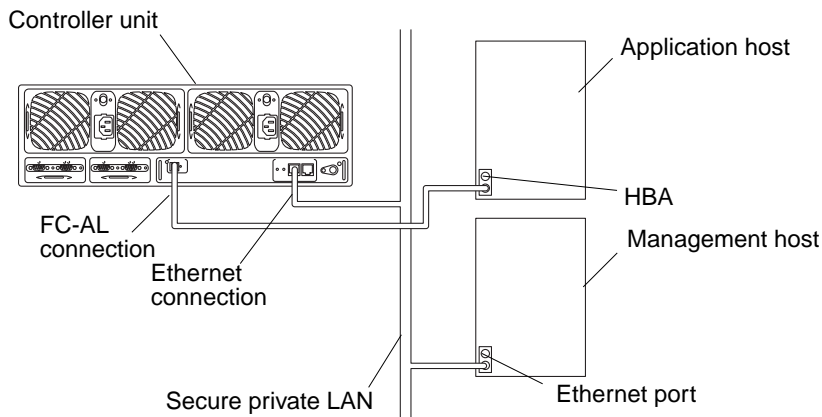


FIGURE 4-1 Single Host Connected to One Controller Unit

Single Host With Two Controller Units Configured as a Partner Group

FIGURE 4-2 shows one application host connected through FC-AL cables to one array partner group, which consists of two Sun StorEdge T3+ arrays. The Ethernet connection from the master controller unit is on a secure private network and requires an IP address for the partner group. In the event of a failover, the alternate master controller unit will use the master controller unit's IP address and MAC address.




FIGURE 4-2 Single Host With Two Controller Units Configured as a Partner Group

This configuration is a recommended enterprise configuration for RAS functionality because there is no single point of failure. For multi-path failover, this configuration supports VERITAS DMP by VERITAS Volume Manager, and the Sun StorEdge Traffic Manager Software (STMS) in the Solaris operating environment. The following three global parameters must be set on the master controller unit:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpzio`
- `sys cache auto`
- `sys mirror auto`

For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Host Multi-Pathing Management Software

While Sun StorEdge T3+ arrays are redundant devices that automatically reconfigure whenever a failure occurs on any internal component, a host-based solution is needed for a redundant data path. Supported multipathing solutions include:

- VERITAS DMP by VERITAS Volume Manager
- Sun StorEdge Traffic Manager software

During normal operation, I/O moves on the host channel connected to the controller that owns the LUNs. This path is a *primary path*. During failover operation, the multipathing software directs all I/O to the alternate channel's controller. This path is the *failover path*.

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

Note – The multipathing software solution must be installed on the application host to achieve a fully redundant configuration.

FIGURE 4-3 shows a failover configuration.




FIGURE 4-3 Failover Configuration

Single Host With Four Controller Units Configured as Two Partner Groups

FIGURE 4-4 shows one application host connected through FC-AL cables to four arrays configured as two separate partner groups. This configuration can be used for capacity and I/O throughput requirements. Host-based multi-pathing software is required for this configuration.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

The following three parameters must be set on the master controller unit:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.




FIGURE 4-4 Single Host With Four Controller Units Configured as Two Partner Groups

Single Host With Eight Controller Units Configured as Four Partner Groups

FIGURE 4-5 shows one application host connected through FC-AL cables to eight Sun StorEdge T3+ arrays, forming four partner groups. This configuration is the maximum allowed in a 72-inch cabinet. This configuration can be used for footprint and I/O throughput.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

The following three parameters must be set on the master controller unit:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Host-based multipathing software is required for this configuration.




FIGURE 4-5 Single Host With Eight Controller Units Configured as Four Partner Groups

Hub Host Connection

The following sample configurations are included in this section:

- “Single Host With Two Hubs and Four Controller Units Configured as Two Partner Groups” on page 35
- “Single Host With Two Hubs and Eight Controller Units Configured as Four Partner Groups” on page 38
- “Dual Hosts With Two Hubs and Four Controller Units” on page 40
- “Dual Hosts With Two Hubs and Eight Controller Units” on page 42
- “Dual Hosts With Two Hubs and Four Controller Units Configured as Two Partner Groups” on page 44
- “Dual Hosts With Two Hubs and Eight Controller Units Configured as Four Partner Groups” on page 46

Single Host With Two Hubs and Four Controller Units Configured as Two Partner Groups

FIGURE 4-6 shows one application host connected through FC-AL cables to two hubs and two array partner groups. The Ethernet connection on the master controller unit is on a secure private network and requires an IP address for the partner group. In the event of a failover, the alternate master controller unit will use the master controller unit’s IP address and MAC address.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting arrays to a hub. Arrays can be connected to any available port on the hub.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator’s Manual* for further details.

The following three parameters must be set on the master controller unit:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Host-based multipathing software is required for this configuration.




FIGURE 4-6 Single Host With Two Hubs and Four Controller Units Configured as Two Partner Groups

Single Host With Two Hubs and Eight Controller Units Configured as Four Partner Groups

FIGURE 4-7 shows one application host connected through FC-AL cables to two hubs and to eight Sun StorEdge T3+ arrays, forming four partner groups. This configuration is the maximum allowed in a 72-inch cabinet. This configuration can be used for footprint and I/O throughput.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting arrays to a hub. An array can be connected to any available port on the hub.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following three parameters must be set on the master controller unit, as follows:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Host-based multipathing software is required for this configuration.




FIGURE 4-7 Single Host With Two Hubs Configured and Eight Controller Units as Four Partner Groups

Dual Hosts With Two Hubs and Four Controller Units

FIGURE 4-8 shows two application hosts connected through FC-AL cables to two hubs and four Sun StorEdge T3+ arrays. This configuration, also known as a *multi-initiator configuration*, can be used for footprint and I/O throughput. The following limitations should be evaluated when proceeding with this configuration:

- Use host-based mirroring software such as VERITAS Volume Manager or Solaris Volume Manager to avoid the risk caused by any array or data path single point of failure.
- When you are configuring more than a single array to share a single FC-AL loop, as with a hub, set the array target addresses to unique values.

This configuration is not a recommended for RAS functionality because the controller is a single point of failure.

Note – There are no hub port position dependencies when connecting arrays to a hub. An array can be connected to any available port on the hub.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following two parameters must be set on the master controller unit, as follows:

- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.




FIGURE 4-8 Dual Hosts With Two Hubs and Four Controller Units

Dual Hosts With Two Hubs and Eight Controller Units

FIGURE 4-9 shows two application hosts connected through FC-AL cables to two hubs and eight Sun StorEdge T3+ arrays. This configuration, also known as a *multi-initiator configuration*, can be used for footprint and I/O throughput. The following limitations should be evaluated when proceeding with this configuration:

- Avoid the risk caused by any array or data path single point of failure using host-based mirroring software such as VERITAS Volume Manager or Solaris Volume Manager.

Note – This configuration, running host-based mirroring features from VERITAS Volume Manager or Solaris Logical Volume Manager, represents four arrays of data mirrored to the other four arrays using host-based mirroring.

- When configuring more than a single array to share a single FC-AL loop, as with a hub, set the array target addresses to unique values.

This configuration is not recommended for RAS functionality because the controller is a single point of failure.

Note – There are no hub port position dependencies when connecting arrays to a hub. An array can be connected to any available port on the hub.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following two parameters must be set on the master controller unit, as follows:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.




FIGURE 4-9 Dual Hosts With Two Hubs and Eight Controller Units

Dual Hosts With Two Hubs and Four Controller Units Configured as Two Partner Groups

FIGURE 4-8 shows two application hosts connected through FC-AL cables to two hubs and four Sun StorEdge T3+ arrays forming two partner groups. This *multi-initiator configuration* can be used for footprint and I/O throughput.

Note – This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting arrays to a hub. An array can be connected to any available port on the hub.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following three parameters must be set on the master controller unit, as follows:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*

Host-based multipathing software is required for this configuration.




FIGURE 4-10 Dual Hosts With Two Hubs and Four Controller Units Configured as Two Partner Groups

Dual Hosts With Two Hubs and Eight Controller Units Configured as Four Partner Groups

FIGURE 4-9 shows two application hosts connected through FC-AL cables to two hubs and eight Sun StorEdge T3+ arrays forming four partner groups. This multi-initiator configuration can be used for footprint and I/O throughput.

This configuration is a recommended enterprise configuration for RAS functionality because the controller is not a single point of failure.

Note – There are no hub port position dependencies when connecting Sun StorEdge T3+ arrays to a hub. An array can be connected to any available port on the hub.

When configuring more than one partner group or a single array to share a single FC-AL loop, as with a hub, array target addresses need to be set to unique values. Assign the array target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to Appendix A of the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following two parameters must be set on the master controller unit, as follows:

- Multi-path support software:
 - DMP: `sys mp_support rw`
 - STMS: `sys mp_support mpxio`
- `sys cache auto`
- `sys mirror auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*

Host-based multipathing software is required for this configuration.




FIGURE 4-11 Dual Hosts With Two Hubs and Eight Controller Units Configured as Four Partner Groups

Switch Host Connection

This section contains the following example configurations:

- “Dual Hosts With Two Switches and Two Controller Units” on page 48
- “Dual Hosts With Two Switches and Eight Controller Units” on page 50

Dual Hosts With Two Switches and Two Controller Units

FIGURE 4-12 shows two application hosts connected through FC-AL cables to two switches and two Sun StorEdge T3+ arrays. This multi-initiator configuration can be used for I/O throughput.

Note – This configuration is not recommended for RAS functionality because the controller is a single point of failure.

Evaluate the following limitations before proceeding with this configuration:

- Avoid the risk caused by any array or data path single point of failure using host-based mirroring software such as VERITAS Volume Manager or Solaris Volume Manager.
- When configuring more than a single array to share a single FC-AL loop, as with a hub, set the array target addresses to unique values.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following two parameters must be set on the master controller unit, as follows:

- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

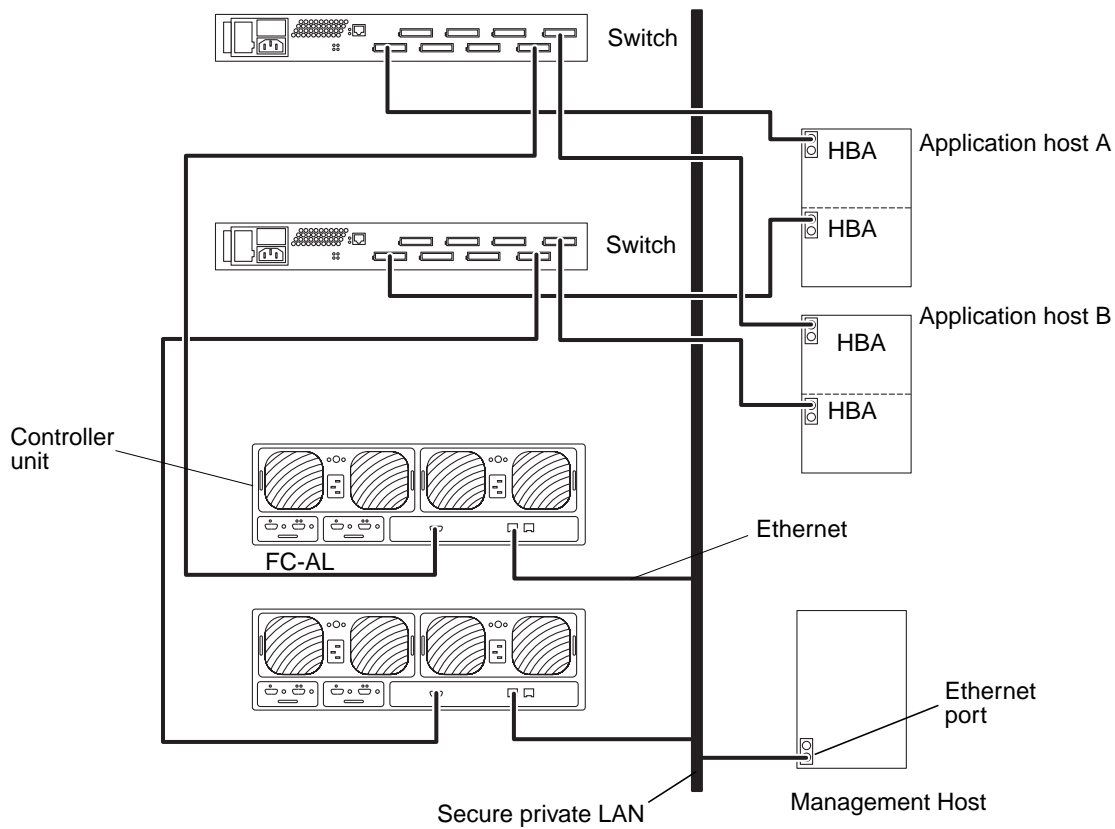


FIGURE 4-12 Dual Hosts With Two Switches and Two Controller Units

Dual Hosts With Two Switches and Eight Controller Units

FIGURE 4-13 shows two application hosts connected through FC-AL cables to two switches and eight Sun StorEdge T3+ arrays. This multi-initiator configuration can be used for footprint and I/O throughput.

Note – This configuration is not recommended for RAS functionality because the controller is a single point of failure.

The following limitations should be evaluated when proceeding with this configuration:

- Avoid the risk caused by any array or data path single point of failure by using host-based mirroring software such as VERITAS Volume Manager or Solaris Logical Volume Manager.
- When configuring more than a single array to share a single FC-AL loop, as with a hub, set the array target addresses to unique values.

Each array needs to be assigned a unique target address using the `port set` command. These target addresses can be any number between 1 and 125. At the factory, the array target addresses are set starting with target address 1 for the bottom array and continuing to the top array. Use the `port list` command to verify that all arrays have a unique target address. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for further details.

The following two parameters must be set on the master controller unit, as follows:

- `cache mode = auto`
- `cache mirroring = auto`

Note – For information on setting these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.




FIGURE 4-13 Dual Hosts With Two Switches and Eight Controller Units

Preparing to Install

This chapter outlines the steps you need to take to prepare your system for installing the product.

Topics in this chapter include:

- “Preparing for the Installation” on page 54
- “To Inspect the Array” on page 55
- “To Edit the Host Files” on page 57
- “To Install Required Patches” on page 59
- “Installation Checklist” on page 61

Preparing for the Installation

Before you install the array, make sure you have completed the following tasks:

- Prepare the area according to the guidelines and procedures in the *Sun StorEdge T3+ Array Site Preparation Guide*.
- Review safety requirements for the array installation in the *Sun StorEdge T3+ Array Regulatory and Safety Compliance Manual*.
- Obtain all the equipment necessary to complete the installation, such as additional cabling or adapters. Refer to the *Sun StorEdge T3+ Array Installation and Configuration Guide* or contact a Sun sales representative to help you determine your configuration needs for additional hardware and software.

The array is designed to be easily installed. A Phillips screwdriver is required for installing the array into a rack; a flat blade screwdriver is helpful for removing and replacing components.

Note – The array configuration requires an Ethernet connection for each controller unit. The Sun StorEdge T3+ array includes an Ethernet cable (category 5) with each array for the 10/100BASE-T network connection.

The array uses nominal input voltages of 100–120 VAC or 200–240 VAC. Sun products are designed to work with single-phase power systems that have a grounded neutral conductor. See Appendix A for additional electrical specifications.



Caution – To reduce the risk of electrical shock, do not connect Sun products into another type of power source. Contact your facilities manager or a qualified electrician if you are unsure what type of power is supplied to your building.

▼ To Inspect the Array



Caution – This procedure requires two people to lift and move the array. Use care to avoid injury. An array can weigh up to 67 pounds (30 kg).

1. Unpack the array.

You should have the following items:

- Sun StorEdge T3+ array
- One 5m fiber-optic cable (unique to the array model)
- Two power cords

The cables are illustrated in FIGURE 5-1.




FIGURE 5-1 Sun StorEdge T3+ Array Cables

2. Inspect the array for evidence of damage.

If the array is damaged, keep all contents and packing materials for the shipping company's agent to inspect.

3. Save the packing materials for future use.

4. Remove the front panel of the array by pressing in on the side latches and pulling the cover off (FIGURE 5-2).




FIGURE 5-2 Removing the Front Panel

- 5. Locate the pull-out tab at the left side of the array next to the first disk drive (FIGURE 5-3).**

This tab contains the array serial number and *media access control (MAC) address*. Record this information; you will need the MAC address for the installation.




FIGURE 5-3 Serial Number and MAC Address on Pull-out Tab

6. Make sure that all the disk drives, the power and cooling units, the interconnect cards, and the controller card (if present) are firmly seated in the unit.
7. Replace the front panel.



Caution – You must replace the front panel for the array to meet regulatory emissions requirements.

▼ To Edit the Host Files

The JumpStart™ feature automatically downloads the assigned IP address to the array. To enable this feature, you must edit your host file before cabling the array. After you cable the array and power on, the IP address is automatically assigned. Before you begin, make sure you have the following information:

- **MAC address.** See FIGURE 5-3 for the MAC address location. See Step 1 of the following procedure for the format of the MAC address.
- **IP address.** For this information, contact the person who maintains your network.
- **Array name.** This is the user-assigned name of the array you are installing.

Note – If you are installing an enterprise configuration (partner group), you need the MAC address that will be assigned only to the master controller unit (also called *master unit*). In a partner group, this is the array that is located on the bottom. In Step 1, enter the MAC address for the master unit. When you install the array in the next section, make sure that the master unit is placed on the bottom of the configuration.

Do *not* use the MAC address of the array that you are placing on the top. In a partner group, the array that is placed on top is referred to as the *alternate master controller unit* or *alternate master unit*.

To set the network IP address for the array:

1. On a host connected to the same subnet as the array, 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 array you are installing.

2. Add the IP address and array name to the `/etc/hosts` file.

For example:

```
192.129.122.111 array-name
```

In this example: 192.129.122.111 is the assigned IP address.

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

To ensure that the Solaris software environment uses the changes made to the `/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.

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

4. Determine if the reverse address resolution protocol (RARP) daemon is running by typing:

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

- If the RARP daemon is running, you must kill the process and restart it for 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, proceed to the next step.

5. Start the RARP daemon in the Solaris software environment by typing:

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

The IP address will automatically download to the array after you install the array and power on.

Note – In some cases, the array could time out before it receives the RARP request through an Ethernet switch. If this happens, the array cannot receive the assigned IP address. If the array should time out before receiving the RARP request, it could be due to an improper spanning-tree setting of the Ethernet switch. Refer to your switch vendor documentation for information on spanning-tree settings and how to change them. Changing this setting properly will enable the array to receive the RARP request before timing out.

▼ To Install Required Patches

Install all the required patches listed below before installing the Sun StorEdge T3+ array. These patches are available on the SunSolveSM web site:

`http://sunsolve.sun.com`

To retrieve patches:

- 1. In the SunSolve Online column, click on Patches.**
- 2. In the Patches column, click on PatchPro.**
- 3. Click on Storage Products.**
- 4. Under Disk Arrays, check the box next to StorEdge T3.**
- 5. Scroll to the bottom of the page and click Generate Patch List.**

Refer to the patch README file for download instructions and patch information.

If you do not find the patch you need in TABLE 5-1:

- a. Backup to the SunSolve online Patches page that was obtained from Step 1.**
- b. Enter a patch ID in the Patch Finder box and click Find Patch.**

Note – For related software products such as VERITAS Volume Manager you might need patches that are not listed in TABLE 5-1. Refer to the SunSolve web site and to the release notes of these products for their required patches.

TABLE 5-1 Required Patches

| System Type | Solaris 2.6 Operating Environment | Solaris 7 Operating Environment | Solaris 8 Operating Environment |
|---|---|---|---|
| All | 105356-18 or later (ssd driver) | 107458-13 or later (ssd driver) | 109524-09 or later (ssd driver) |
| | 106226-02 or later (format patch) | 107473-07 or later (luxadm patch) | |
| | 105181-29 or later (kernel update patch) | 107834-03 or later (DKIO Extensions Patch) | |
| | | 106541-18 or later (Kernel Update Patch) | |
| VERITAS VM 3.1 | 110253-04 | 110254-04 | 110255-04 |
| VERITAS VM 3.1.1 | 110451-05 | 110452-05 | 111118-05 |
| VERITAS VM 3.2 | 111907-04 | 111908-04 | 111909-04 |
| Volume Manager Storage Administrator | 111904-04 | 111904-04 | 111904-04 |
| PCI (ifp) Only | 107280-08 or later (ifp/PCI driver/PCI Systems only) | 107292-08 or later (ifp/PCI driver/PCI Systems only) | 109189-03 or later (ifp/PCI driver/PCI Systems only) |
| | 109399-03 or later (PCI Host Adapter Firmware Fcode) | 109399-03 or later (PCI Host Adapter Firmware Fcode) | 109399-03 or later (PCI Host Adapter Firmware Fcode) |
| SBus/sf-socal Only | 105375-26 or later (sf/socal driver/SBus systems only) | 107469-08 or later (sf/socal driver/SBus systems only) | 109460-06 or later (sf/socal driver/SBus systems only) |
| | 109400-03 or later (SBus Host Adapter Firmware Fcode) | 109400-03 or later (SBus Host Adapter Firmware Fcode) | 109400-03 or later (SBus Host Adapter Firmware Fcode) |
| Sun StorEdge Network Foundation Software on PCI Network Adapters¹ | | | 111095-04 111097-04 111412-04 111413-04 |

1. Applies only to the following PCI network adapters: Sun StorEdge PCI Single Fibre Channel network adapter, Sun StorEdge PCI Dual Fibre Channel network adapter, and Sun StorEdge CompactPCI Dual Fibre Channel network adapter.

Installation Checklist

- Diagram of lab and network configuration
- Host ID
- Host name
- Host IP address
- Gateway IP address
- Gateway netmask
- Array IP address
 - Top array IP address
 - bottom array IP adders
- Array name
 - Top array name
 - Bottom array name
- root password for array
- TFTP IP host address
- OS/patch revision level
- Veritas DMP release
- Primary application
- StarTool release
- Location of firmware files
- RAID controller firmware file names
- EPROM firmware file name
- UID firmware file name
- Disk drive firmware file name

Array Installation

This chapter describes how to install the Sun StorEdge T3+ array in either a workgroup (single) or enterprise (partner group) configuration.

Note – This chapter is for installing new array units only. To reconfigure existing single array units into an enterprise configuration, contact your SunService representative.



Caution – In an enterprise configuration, interconnected arrays must be of the same model type; that is, you can only connect two Sun StorEdge T3+ arrays together. A mixed array partner group is not supported.



Caution – For security reasons, connect the array to a private network only.

The following sections describe specific steps of the installation.

Note – Follow these directions in the order provided to ensure a successful installation.

Topics in this chapter include:

- “To Install the Array” on page 64
- “To Connect the Cables” on page 65
- “To Power On and Verify the Hardware Configuration” on page 70
- “To Establish a Network Connection” on page 72
- “To Verify the Firmware Level and Configuration” on page 75
- “Defining and Mounting Volumes” on page 81
- “Changing the Default Configuration” on page 88
- “Connecting the Host System” on page 89
- “Powering Off and On” on page 94

▼ To Install the Array

This section describes the procedures for installing either a workgroup or enterprise configuration; see Chapter 1 for an explanation of each configuration. The array can either be placed on a tabletop or rackmounted in a server cabinet or expansion cabinet. This section provides guidelines for tabletop placement of an array configuration. To rackmount the array in a cabinet, you must order a array rackmount kit. Instructions for rackmounting arrays are included in the kit.

The array is designed to sit on a desk or table. Use the following guidelines to prepare a location for your system.

- Choose a desk or table that can support up to 67 pounds (30 kg) for one fully configured array or 135 pounds (60 kg) for two arrays.
- Leave enough space in front and in back of the array to access components.
- Provide a minimum space of 6 inches (15 cm) in front and in back of the array for adequate air flow.
- Keep power and interface cables clear of foot traffic. Route cables inside walls, under the floor, through the ceiling, or in protective channels. Route interface cables (excluding fiber-optic cables) away from motors and other sources of magnetic or radio frequency interference.
- Make sure that the fiber-optic and power cable lengths do not exceed cable length limitations. See Appendix A for maximum cable lengths.
- Ensure that the operating environment for the array does not exceed the specifications. See Appendix A for environmental specifications.



Caution – Use two people to lift the array to avoid injury. It can weigh up to 67 pounds (30 kg).

1. Place the array horizontally in the designated location.



Caution – Do not place the array in a vertical position. It is not designed to stand vertically.

2. If you are installing a second array, stack it on top of the first array.

The second array is the *alternate master unit*. The first array, or the bottom array in the stack, is the *master unit*.



Caution – Do not stack more than two arrays together. Stacking more than two arrays could damage them.

You are now ready to connect the cables.

▼ To Connect the Cables

This section explains how to connect cables for both a workgroup (standalone) and for an enterprise (partner group) configuration.

Note – If you are connecting a single array, skip the steps that start with “Partner group only” and proceed to the next step.

Before you begin, make sure you have all of the required cables and adapters:

- Fiber-optic cable, one per unit
- AC power cords, two per unit
- Interconnect cables, one pair of cables per partner group
- Ethernet interface cable, 1 per controller unit:
 - 10BASE-T (category 3) interface cable for the array (must be purchased separately)
 - 10/100BASE-T (category 5) interface cable for the Sun StorEdge T3+ array (included with array packaging)

Sun StorEdge T3+ array cables are shown in FIGURE 6-1

Note – Your array packaging might not have included interconnect cables. If you want to install an enterprise configuration and do not have interconnect cables, contact your Sun sales representative.




FIGURE 6-1 Sun StorEdge T3+ Array Cables

- 1. Connect the fiber-optic cable to the host bus adapter (HBA).**
Make sure you attach the standard connector (SC) on the fiber-optic cable to the HBA.
- 2. Connect the other end of the fiber-optic cable to the FC-AL connector on the controller card at the back of the master (bottom) unit.**
Connect the LC-SFF (small form factor) end of the cable to the FC-AL interface (FIGURE 6-2).
- 3. Partner group only: Repeat Step 1 and Step 2 to connect a second fiber-optic cable to the alternate master (top) unit and to a separate HBA.**




FIGURE 6-2 Connecting the Fiber-Optic Cable to the FC-AL Connector on a Sun StorEdge T3+ Array

- 4. Attach one end of the Ethernet cable to the Ethernet port on the controller card of the master unit (FIGURE 6-3).**

Use a category 5 Ethernet cable.




FIGURE 6-3 Connecting the Ethernet Cable

5. **Attach the other end of the Ethernet cable to a network hub or router.**

Note – You must use a shielded Ethernet interface cable to comply with regulatory emissions requirements.

6. **Partner group only: Repeat Step 4 and Step 5 to connect the alternate master unit to a hub or router using a second Ethernet cable.**

Connect this cable to a hub or router on the same network as the master unit.

7. **Attach the two power cords to the two power and cooling units on the back of the master unit (FIGURE 6-4).**




FIGURE 6-4 Connecting the Power Cords

8. **Partner group only: Repeat Step 7 to connect the power cords to the alternate master unit.**



Caution – Do *not* power on the units yet. If you are connecting a single array unit, proceed to “To Power On and Verify the Hardware Configuration” on page 70 for instructions on powering on the unit. If you are connecting a partner group, proceed to Step 9.

9. Partner group only: Connect the interconnect cables to the interconnect cards as shown in FIGURE 6-5.

Make sure you connect the cables to the correct interconnect card connectors exactly as shown in the figure. This cable connection determines the master and alternate master relationship. Tighten the retaining screws.




FIGURE 6-5 Connecting the Interconnect Cables

A fully cabled partner group is shown in FIGURE 6-6.




FIGURE 6-6 Fully Cabled Partner Group (Enterprise Configuration)

▼ To Power On and Verify the Hardware Configuration

To power on and verify the hardware configuration:

- 1. Connect the power cords from each array into AC outlets.**
Use alternate power sources to make sure that power and cooling unit features are redundant.
- 2. Verify that AC power is present on each power and cooling unit.**
When AC power is present, the AC LED on each power and cooling unit is solid amber and the fans turn at low speed.
- 3. Press the power button on each power and cooling unit on all arrays to power on the units.**
When the arrays are powered on, the AC LED and power supply (PS) LED on the power and cooling units are green.




FIGURE 6-7 Power and Cooling Unit LEDs and Power Button Location

4. Check the LEDs at the front and back of the each array to ensure that all components are receiving power and are functional.

While the drives are spinning up, the LEDs will blink. The array boot time might take several minutes, after which all LEDs should be solid green, indicating the unit is receiving power and there is no drive activity.

For more information on the LEDs and how to interpret them, see “Array LEDs” on page 156.

Note – The batteries in the power and cooling units recharge after you power on the unit. While the batteries are recharging, write-behind cache is disabled.

▼ To Establish a Network Connection

Use the Ethernet network connection to initialize the network addresses and to administer the array. The IP address was automatically downloaded by the JumpStart program when you powered on the array.

Note – Have the gateway IP address information before beginning this procedure. Contact the person who maintains your network for this information.

Note – The telnet session will not connect until the controller LEDs are green.

1. **On the host, use the `telnet` command with the array name (or IP address) to connect to the array.**

For example:

```
# telnet array-name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

pSOSystem (129.150.47.101)
```

Note – The telnet session verifies whether your network connection is good. If you cannot connect through the telnet session, see the troubleshooting information in the *Sun StorEdge T3+ Array Administrator's Manual* to analyze the problem.

2. **Type `root` as the login and press Return at the password prompt.**

The array displays the banner information, which appears similar to the following:

```
Login: root
Password: <Return>

T3B Release 2.00 2001/04/02 15:21:29 (192.168.209.243)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.

:/:<1>
```

3. Use the `passwd` command to set a password for the root account.

Press Return when prompted for the OLD password.

```
:/:<1>passwd
OLD password: <Return>
NEW password: New password
NEW password (confirm): New password
:/:<2>
```

Note – It is important to set a root password for security reasons.

4. Set the gateway using the `set gateway` command.

The gateway address enables you to access the array outside the subnet.

For example:

```
:/:<3> set gateway 129.150.47.1
```

5. Set the netmask using the `set netmask` command.

The netmask specifies the network mask used to implement IP subnetting.

```
:/:<4> set netmask 255.255.255.0
```

6. Set the host name using the `set hostname` command.

```
:/:<5> set hostname hostname
```

7. Use the `set timezone` command to set the time zone region.

For the time zone region, see the *Sun StorEdge T3+ Array Administrator's Manual*. For example, if you are located in the Pacific Standard Time (PST) time zone, the region would be `America/Los_Angelos` as shown:

```
:/:<6> set timezone :America/Los_Angelos
```

8. Use the `set` command to confirm the setting.

```
:/:<7> set
bootmode      tftp
bootdelay     3
sn            103676
ip            10.4.31.28
netmask       255.255.255.0
gateway       10.4.31.1
tftphost      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
```

9. Set the date using the `date` command.

The date syntax is `yyyymmddHHMM.SS`.

For example:

```
:/:<9> date 200005171620.30
Wed May 17 16:20:30 2000
:/:<10>
```

10. Reset the array using the `reset` command.

Answer **Y** when prompted for confirmation. A reset disconnects the telnet session with the array.

```
:/:<110> reset
Reset the system, are you sure? [N]: Y
```

The array reboots. This can take up to several minutes.

11. On the host, exit the telnet session by pressing Control and the right bracket (]) symbol.

The telnet prompt is displayed. Type `quit`.

```
# <Control>]
telnet>quit
#
```

▼ To Verify the Firmware Level and Configuration

The array contains various types of firmware that are peculiar to the array model.

The Sun StorEdge T3+ array contains the following firmware:

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

Use this procedure to make sure you are running the latest firmware versions on your particular array model, and that the array configuration information indicates that the unit is ready for operation. You can check the firmware versions and array information in a telnet session with the array.

Note – Verification of level 1 firmware on a Sun StorEdge T3+ array is not required at this time. Access to level 1 firmware operations is limited to qualified service providers.

1. On the host, use the `telnet` command with the *array name* (or IP address) to connect to the array.

For example:

```
# telnet array-name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

pSOSystem (129.150.47.101)
```

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

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

- Here is 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 2.00 level 3 controller firmware.

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

For example (enterprise configuration):

```

:/:<6>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 |

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 level 2 controller firmware, the level 2 image is 0200.
5. **Refer to the SunSolve web site at <http://sunsolve.sun.com> for the latest firmware release information:**
 - a. **In the SunSolve Online column, click on Patches.**
 - b. **In the Patches column, click on PatchPro.**
 - c. **Click on Storage Products.**

Refer to the patch README file for details on firmware and specific firmware patch requirements.

Note – If you do not have access to the SunSolve web site, contact your SunService provider for patch information.

6. **If you need to upgrade the firmware, do so now by following the instructions in Chapter 7.**

If you do not need to upgrade firmware, proceed to Step 7.

7. **Type `port list` to display how the paths are mapped from the host to the volumes.**

An example of `port list` output for a single controller unit:

```

:/:<7>port list

port  targetid  addr_type  status  host  wwn
ulp1  1           hard      online  sun   50020f23000058de

```

Use the world wide name (`wwn`) from the `port list` output to map the volume to the Solaris environment using the `format(1M)` utility on the data host. In the following example, the output for the `format` command shows the same `wwn`, listed in the last line, as the `wwn` in the `port list` output.

```

# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
    /pci@1f,4000/scsi@3/sd@0,0
  1. c2t1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
    /pci@6,2000/SUNW,ifp@1/ssd@w50020f23000002ba,0

```

8. Type `sys stat` to view the controller status.

- An example of `sys stat` output for a single controller unit:

```
:/:<9>sys stat
Unit   State   Role   Partner
-----
 1     ONLINE Master
```

- An example of `sys stat` output for a partner group:

```
:/:<10>sys stat
Unit   State   Role   Partner
-----
 1     ONLINE Master   2
 2     ONLINE AlterM  1
```

9. Type `vol mode` to display the cache and mirror settings and to determine whether the volume is mounted.

The default settings for a single controller unit are as follows:

```
:/:<11>vol mode
volume      mounted  cache      mirror
v0          yes      writebehind off
```

10. Type `fru stat` to display the health status of each FRU.

All FRUs must be in optimal condition, as shown in this example for a workgroup configuration:

```

:/:<12>fru stat
CTLR      STATUS  STATE   ROLE     PARTNER  TEMP
-----  -
ulctr     ready   enabled master    -         29.0

DISK      STATUS  STATE   ROLE     PORT1    PORT2    TEMP  VOLUME
-----  -
uld1     ready   enabled data disk ready    ready    30    v0
uld2     ready   enabled data disk ready    ready    30    v0
uld3     ready   enabled data disk ready    ready    34    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    36    v0
uld7     ready   enabled data disk ready    ready    37    v0
uld8     ready   enabled data disk ready    ready    41    v0
uld9     ready   enabled standby ready     ready    34    v0

LOOP      STATUS  STATE   MODE     CABLE1   CABLE2   TEMP
-----  -
ull1     ready   enabled master    -         -         29.5
ull2     ready   enabled slave    -         -         30.0

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
    
```

Defining and Mounting Volumes

This section includes information on defining and mounting volumes for both a workgroup and enterprise configuration. For information on the commands used in this section, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

The array has been set up with a default RAID 5 configuration. To view the configuration, use the `vol list` command while in a Telnet session with the array. For example:

```
:/:<1>vol list

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

Where:

- v0 is the volume.
- 143.2 GB is the total disk space of the volume.
- 5 is the RAID level.
- u1d1-9 is the number of the array (u1) and the numbers of the disk drives (d1 through d9).
- none indicates that there is no standby (hot spare) disk drive.

The physical location of the disk drives is shown in FIGURE 6-8.




FIGURE 6-8 Disk Drive Physical Numbering

- If you want to change the default configuration, do not continue with this section. Instead, refer to the *Sun StorEdge T3+ Array Administrator's Manual* for information on how to remove the default volume and reconfigure the disk drives.
- If you have a *workgroup* configuration and want to mount the default volume, proceed to “To Configure the Workgroup” on page 83.
- If you have an *enterprise* configuration and want to mount the default volumes, skip to “To Configure the Enterprise” on page 85.

Note – If you are creating new volumes or changing the volume configuration, you must first manually rewrite the label of the previous volume using the `autoconfigure` option of the `format(1M)` utility. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more information on this procedure.

Note – The default configuration does not include a hot spare. If you want to add a hot spare to the configuration, you must remove the existing volume and re-create the configuration. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for information on adding a hot spare.

▼ To Configure the Workgroup

To display and mount the default volume of a single controller unit:

1. On the array, type `vol stat` to check the status of the volume.

```
:/:<2> vol stat
v0          u1d1  u1d2  u1d3  u1d4  u1d5  u1d6  u1d7  u1d8  u1d9
unmounted   0      0      0      0      0      0      0      0      0
```

Where:

- `v0` is the name of the volume.
- `unmounted` is the status of the volume.
- `u1d1-u1d9` are the disk drives in the volume.
- `0` is the status of the disk drives in the volume.

All the disk drives should show a status of 0. Other possible status messages are listed in the following table.

TABLE 6-1 Drive Status Messages

| Value | Description |
|-------|--|
| 0 | Drive is mounted |
| 2 | Drive is present |
| 3 | Drive is spun up |
| 4 | Drive is disabled |
| 5 | Drive has been replaced |
| 7 | Invalid system area exists on drive |
| 9 | Drive is not present |
| D | Drive is disabled and is being reconstructed |
| S | Drive is substituted |

- If the `vol stat` output shows the volume as *mounted*, skip to Step 4.
- If the `vol stat` output shows the volume as *unmounted*, proceed to Step 2.

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

```
:/:<3>vol mount volume-name
```

3. Use the `vol stat` command to verify that the volume is mounted.

```
:/:<5> vol stat
v0          u1d1    u1d2    u1d3    u1d4    u1d5    u1d6    u1d7    u1d8    u1d9
mounted     0        0        0        0        0        0        0        0        0
```

After the volume is mounted, it becomes available to the attached host system.

Note – If you are running on the Solaris 7, 11/99 operating environment or a later release of the Solaris operating environment, skip Step 4 and proceed to Step 5. Later versions of the Solaris operating environment automatically recognize added storage devices without additional command operations.

4. Use the `luxadm(1M)` command to recognize the new volume.

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

a. Make sure there is a `/dev/es` directory on the host system. If not, type:

```
# mkdir /dev/es
```

The `/dev/es` directory is necessary for running the `luxadm` command.

b. On the host system, type `luxadm insert`:

```
# luxadm insert
```

Note – If the `luxadm` utility is not available, you must do a reconfiguration reboot (`boot -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 `boot -r` command.

5. Verify that the new volume is available to the host system by using the `format(1M)` command on the data host.

```
7# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
    /pci@1f,4000/scsi@3/sd@0,0
  1. c2t1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
    /pci@6,2000/SUNW,ifp@1/ssd@w50020f23000002ba,0
```

Proceed to “Connecting the Host System” on page 89 to continue the installation.

▼ To Configure the Enterprise

To define and mount the default volumes of an enterprise configuration, you must first make sure the volume on the master unit is mounted, and then add a volume on the alternate master as follows:

1. On the array, type `vol stat` to check the status of the volume.

Only the volume on the master controller unit is displayed.

```
:/:<1>vol stat
v0          u1d1    u1d2    u1d3    u1d4    u1d5    u1d6    u1d7    u1d8    u1d9
unmounted   0        0        0        0        0        0        0        0        0
```

Where:

- `v0` is the name of the volume.
- `unmounted` is the status of the volume.
- `u1d1-u1d9` are the disk drives in the volume.
- `0` is the status of the disk drives in the volume.

All the disk drives should show a status of 0. Possible status messages are listed in TABLE 6-1.

- If the `vol stat` output shows the volume as *mounted*, skip to Step 4.
- If the `vol stat` output shows the volume as *unmounted*, proceed to Step 2.

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

```
:/:<2>vol mount v0
```

3. Type `vol stat` to verify that the volume on the master controller unit is mounted.

```
:/:<3> vol stat
v0          u1d1  u1d2  u1d3  u1d4  u1d5  u1d6  u1d7  u1d8  u1d9
mounted    0      0      0      0      0      0      0      0      0
```

4. Use the `vol add` command to create the default volume on the alternate master as follows:

- a. Define the volume name (`vol add v1`).
- b. Define the drives where the volume will reside (`data u2d1-9`).
- c. Define the RAID level (`raid 5`).

```
:/:<4>vol add v1 data u2d1-9 raid 5
```

Where:

- `v1` is the volume name.
- `u2d1-9` indicates the location of the volume: unit 2, disk drives 1 through 9.
- `raid 5` is RAID level 5.

5. Type `vol stat` to check the status of the volumes.

The status of all drives must be 0. For example:

```
:/:<5>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
unmounted  0      0      0      0      0      0      0      0      0
```

6. Use the `vol init` command to initialize the volume on the alternate master.

```
:/:<6>vol init v1 data
```

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.

7. Use the `vol mount` command to mount the volume on the alternate master.

```
:/:<7>vol mount v1
```

8. Type `vol list` to confirm that you created the volume correctly.

```
:/:<8>vol list

volume          capacity  raid  data      standby
v0              143.2 GB   5     u1d1-9    none
v1              143.2 GB   5     u2d1-9    none
```

Note – If you are running on the Solaris 7, 11/99 operating environment or a later release of the Solaris operating environment, skip Step 9 and proceed to Step 10. Later versions of the Solaris operating environment automatically recognize added storage devices without additional command operations.

9. Use the `luxadm(1M)` command to recognize the new volume.

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

a. Make sure there is a `/dev/es` directory on the host system. If not, type:

```
# mkdir /dev/es
```

The `/dev/es` directory is necessary for running the `luxadm` command.

b. On the host system, type `luxadm insert`:

```
# luxadm insert
```

Note – If the `luxadm` utility is not available, you must do a reconfiguration reboot (`boot -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 `boot -r` command.

10. Verify that the new volume is available to the host system by using the `format(1M)` command on the data host.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
    /pci@1f,4000/scsi@3/sd@0,0
  1. c2t1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
    /pci@6,2000/SUNW,ifp@1/ssd@w50020f23000002ba,0
  2. c1t1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
    /pci@6,0/SUNW,socal@1,0/sf@0,0/ssd@w50020f2300000121,0
```

Changing the Default Configuration

The default configuration for the Sun StorEdge T3+ Array is a RAID 5 volume.

- If you want to change the default configuration, do so now before continuing the installation process. Configuration options for the array are discussed in Chapter 8 and Chapter 9. The command-line procedures for reconfiguring the array are in the *Sun StorEdge T3+ Array Administrator's Manual*. Once you have reconfigured the array, return to this chapter and proceed to “Connecting the Host System” on page 89.
- If you are *not* changing the default configuration, proceed to “Connecting the Host System” on page 89.

Connecting the Host System

This section contains procedures that are performed on the host system for the array. Have the documentation for the host available to assist you in completing these tasks.

- “Verifying the Data Host Connection” on page 89
- “Establishing Logical Volumes on the Host” on page 90
- “Creating Soft Partitions on the Data Host” on page 90
- “To Establish Alternate Pathing on the Host” on page 91

Verifying the Data Host Connection

After the physical connection between the host and the array has been established, you can use a utility, such as the `format(1M)` command in the Solaris environment, to verify the existence of a volume on the array. For example:

```
# format
Searching for disks...done

clt1d0: configured with capacity of 133.38GB

AVAILABLE DISK SELECTIONS:
  0. c0t2d0 <drive type unknown>
     /sbus@1f,0/SUNW,fas@e,8800000/sd@2,0
  1. c0t3d0 <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
     /sbus@1f,0/SUNW,fas@e,8800000/sd@3,0
  2. clt1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
     /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w50020f2300000121,0
Specify disk (enter its number):
```

In this example, device number 2 is the array, as identified by the SUN-T300-0113 (0113 changes) label.

Note – If you are creating new volumes or changing the volume configuration, you must first manually rewrite the label of the previous volume using the `autoconfigure` option of the `format(1M)` utility. Refer to the *Sun StorEdge T3+ Array Administrator’s Manual* for more information on this procedure.

Establishing Logical Volumes on the Host

Using the `format(1M)` utility within the Solaris environment, you can view the disk space available on the array from the host. At this point in the installation process, you can use this disk space as you would any physical disk. This includes performing operations such as installing a file system, using the device as a raw device, or partitioning the device. Refer to the data host documentation for more information on establishing logical volumes.

Note – Two commands are available to determine the drive volume capacity of the array. Within the array, use the `vol` command to display the raw volume capacity. Within the Solaris environment, use the `format(1M)` utility (among others) to display the capacity of a newly created volume. The volume capacities differ between the internal array `vol` command and the Solaris environment `format(1M)` command; the Solaris `format(1M)` command reports a smaller volume capacity. This is due to the Solaris environment reporting volume sizes using 1,024 bytes per Kbyte capacities.

For example, using the array `vol` command, a 4-drive, 18-Gbyte RAID 0 volume would report a 71.6-Gbyte capacity within the array. This is calculated as follows:

$$139857920 \text{ blks} * 512 \text{ bytes/blk} * 1 \text{ Gbyte}/1000^3 = 71.6 \text{ Gbytes}$$

The `format(1M)` command in the Solaris environment reports a 66.69 Gbyte capacity for the same volume using 1,024 bytes per Kbyte, calculated as follows:

$$139857920 \text{ blks} * 512 \text{ bytes/blk} * 1 \text{ Gbyte}/1024^3 = 66.69 \text{ Gbytes}$$

There is no difference in the actual disk capacity.

Creating Soft Partitions on the Data Host

The array's native volume management can support a maximum of two volumes or *logical unit numbers* (LUNs) per array unit. This can result in very large volumes (128-Gbytes in a configuration of single 7+1 RAID 5 LUN plus hot spare, with 18-Gbyte drives). Some applications cannot use such large volumes effectively. Two solutions can be used separately or in combination.

- First, use the partitioning utility available on the data host's operating system. In the Solaris environment, this is the `format` utility, which can create up to seven distinct partitions per volume. Note that in the case of the configuration described above, if each partition is equal in size, this will result in 18-Gbyte partitions, which still may be too large to be used efficiently by legacy applications.

- Second, you can use third-party software on the host to create as many partitions as desired from a given volume. In the Solaris environment, VERITAS Volume Manager can be used for this purpose.

For information on using the `format` utility, refer to the `format(1M)` man page. For more information on third-party software such as VERITAS Volume Manager, refer to the documentation for that product.

Note – This completes the installation process for an array workgroup installation. If you are installing an enterprise configuration, you must continue to the next section.

▼ To Establish Alternate Pathing on the Host

This task applies only when two controller units are configured as a partner group in an enterprise configuration.

In a partner group, controller caches are mirrored so that if one controller fails, no data is lost. Any data that is in the cache of the failed controller, but is not yet written to disk, is preserved in the cache of the partner controller and is written to disk by means of one of the back-end Fibre Channel loops connecting the two units.

The back-end Fibre Channel loops are used to maintain host access to the disks of the unit with the failed controller (or any failure in the host data path to that controller). This requires the use of a host-based alternate pathing mechanism such as the Dynamic Multi-Pathing (DMP) feature of VERITAS Volume Manager, Sun Enterprise Alternate Pathing software, or Sun StorEdge Traffic Manager software. When a failure occurs in the host data path, the alternate pathing mechanism reroutes data to the other array controller in the partner group. The controller then redirects the data across one of the back-end loops to the disks of the other controller, thus maintaining data availability to the host application.

For more information and for instructions on installing and configuring alternate pathing software, refer to the documentation for the alternate pathing tool selected.

For the alternate pathing software to function properly, you must enable multipathing support on the array for enterprise configurations. To do so, use the following procedure:

- 1. If you are not already in a Telnet session with the array, start one.**

If you are in a Telnet session, skip to Step 2.

- a. On the host, use the `telnet` command with the array name (or IP address) to connect to the array.

```
# telnet array-name
Trying 129.150.47.101...
Connected to 129.150.47.101.
Escape character is '^]'.

pSOSystem (129.150.47.101)
```

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

2. At the array prompt, enable the alternate pathing software as follows:

- For hosts running Sun StorEdge Traffic Manager Alternate Pathing software, type:

```
:/:<1>sys mp_support mpxio
```

- For hosts running other alternate pathing software, type:

```
:/:<1>sys mp_support rw
```

3. Use the `sys list` command to verify that the `mirror` setting is set to `auto`.

```
:/:<2>sys list
blocksize      : 64k
cache          : auto
mirror         : auto
mp_support     : rw <or mpxio>
naca           : off
rd_ahead      : on
recon_rate     : med
sys memsize    : 128 MBytes
cache memsize  : 1024 MBytes
```

Note – For data hosts running Solaris operating environments earlier than Solaris 7, 11/99 or later, run the `luxadm insert` command so the host can recognize the new paths.

4. Type `port list` to display the world-wide name (wwn).

An example of `port list` output for a partner group:

```
:/:<7>port list

port   targetid  addr_type  status  host  wwn
ulp1   1          hard       online  sun   50020f23000003c3
u2p1   2          hard       online  sun   50020f2300000341
```

5. Use the `wwn` from the `port list` output to correlate the volume to the Solaris environment by using the `format(1M)` utility on the data host.

In the following example, the output for the `format` command would show the same `wwn`, listed in the second to last line (`u1`) and last line (`u2`), as the `wwn` in the `port list` output.

```
format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
     /pci@1f,4000/scsi@3/sd@0,0
  1. c1t0d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
     /pci@6,2000/SUNW,ifp@1/ssd@w50020f23000003c3,0
  2. c2t1d0 <SUN-T300-0113 cyl 34145 alt 2 hd 64 sec 128>
     /pci@6,2000/SUNW,ifp@1/ssd@w50020f2300000341,0
```

6. On the host, exit the Telnet session by pressing Control and the right bracket (]) symbol (Solaris operating environment only) or using the equivalent command for your operating system.

This command takes you to the telnet command prompt; type `quit`.

Note – After multipathing support is enabled and the host is rebooted, from the host side you will see two paths available to the same volume: a primary path and an alternate path. This means that if you are using the `format(1M)` utility in the Solaris environment, you will see two paths for each volume.

This completes the installation process for an enterprise configuration. For information on available array settings and monitoring options, see Chapter 8 and Chapter 9.

Powering Off and On

You do not need to power off the array except under the following circumstances:

- Hardware reconfiguration
- Relocation
- Centerplane and chassis replacement
- Emergency, such as a threat of flooding

▼ To Power Off



Caution – This procedure will power off the array and all arrays connected to it.

Note – If the array is operating in the Solaris environment and there are volumes mounted, unmount the volumes on the data host before powering off the array.

1. **Ensure that the data path between the host and the array has been quiesced.**

There must not be any I/O activity.

2. **In a telnet session with the array, type:**

```
:/:<1>shutdown  
Shutdown the system, are you sure? [N]: y
```

3. **Press the power button once on each power and cooling unit to turn the switch off (FIGURE 6-9).**




FIGURE 6-9 Power Switch Locations

▼ To Power On

When powering on a configuration, always power on the arrays before booting the host server.

- **Power on the master unit by pressing the power button on each power and cooling unit.**

At this point, all power and cooling unit LEDs on the array and any units connected to it will turn green, indicating that power has been restored.

Note – If the green power and cooling unit LEDs on connected units do not light, press the power switches on those units.

Controller Firmware Version 2.1 Upgrade

This chapter explains how to upgrade the Sun StorEdge T3+ array controller firmware version 2.0 or 2.0.1 to version 2.1.

This chapter contains the following topics:

- “Upgrading Overview” on page 98
- “Upgrading Requirements” on page 98
- “Upgrading Limitations” on page 99
- “Connecting to the Array” on page 100
- “To Determine the Firmware Version” on page 103
- “Upgrading the System Firmware” on page 104
- “To Upgrade the Controller Firmware” on page 109
- “To Verify the Upgrade” on page 111

Upgrading Overview

To upgrade controller firmware version 2.1, perform the following steps:

1. **Check to see that the requirements are met. See “Upgrading Requirements” on page 98.**
2. **Make sure that your requirements are not in conflict with the section “Upgrading Limitations” on page 99.**
3. **Read the section “Connecting to the Array” on page 100.**
4. **Determine your controller firmware version to verify if an upgrade is needed. See “To Determine the Firmware Version” on page 103.**
5. **Run the script to automatically install the firmware patches. See “To Install the System Firmware 2.1 Patch” on page 104.**
6. **Upgrade the controller firmware. See “To Upgrade the Controller Firmware” on page 109.**
7. **Check the system functionality. See “To Verify the Upgrade” on page 111.**

Upgrading Requirements

1. This procedure only works on a T3+ array with firmware version 2.0 or 2.0.1.
2. The Sun StorEdge T3+ array(s) must be offline. Stop all processes, applications, and daemons.
3. The contents of patch ID# 112276-02 and controller firmware version 2.1 must be resident on the management host. The software is located on the SunSolve web site:
`http://sunsolve.sun.com`
4. You need to know the Sun StorEdge T3+ array root (supervisor) password.

Upgrading Limitations

1. Once the upgrade is completed and volume slicing is enabled, you cannot revert back to the prior controller firmware version without reconfiguring the entire system. This would require that you:
 - a. Back up all the data to a safe location.
 - b. Reinstall the earlier version controller firmware.
 - c. Restore the data.
2. During the upgrade, entire volumes are mapped (1:1) to slices (regardless of the amount of data in the volume). You will not be able to further slice the volume. However, if you do decide to further slice the volume(s), the following steps must be followed:
 - a. Back up all the data that is on the volume created before the upgrade.
 - b. Remove any existing slices created during the upgrade.
 - c. Create new slices as required.
 - d. Restore the data to a newly created slice that has the capacity to hold the required amount of data.

Connecting to the Array

This section contains the following topics:

- “To Establish an ftp Session” on page 100
- “To Establish a Telnet Session” on page 101

▼ To Establish an ftp Session

To establish an ftp session:

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

For example:

```
mgmt-host: /:<15>ftp 123.123.123.2
Connected to 123.123.123.2.
Escape character is '^]'.

123.123.123.2 pSOSystem (NUPPC/2.0.0-G) ready

Name (123.123.123.2:root):
```

2. **Log in to the array by typing `root`.**

```
Name (123.123.123.2:root): root

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

where *password* is the root password.

Note – Be sure to set the Binary mode if you are transferring firmware.

▼ To Establish a Telnet Session

The Telnet session is a direct network link to the array unit through the command-line interface (CLI). You can execute individual commands to query and repair the unit from this interface. The Telnet session requires access to the unit's IP address. The CLI can then be run from any host that is part of the same subnetwork.

The advantages that a Telnet connection provides over a serial port connection are as follows:

- You can have multiple windows open for each array.
- The Telnet connection provides a faster interface than the serial port connection, which can be useful for displaying `syslog` information.

You can quickly determine the status of the array unit from the CLI. The `syslog` file of the array file system contains a record of events that have occurred in the unit and can also be examined through the CLI.

To open a Telnet connection and start a session with the array:

1. **On the host, use the `telnet` command with the array name (or IP address) to connect to the array.**

For example, to use the `telnet` command to a array named T3-1:

```
mgmt-host# telnet T3-1
> Trying 123.123.123.1...
> Connected to T3-1.
> Escape character is '^]'.

> pSOSystem (123.123.123.1)

> Login: root
> Password: password
```

where *password* is the root password.

2. **Verify the array has a root password by typing it at the prompt.**

If no root password is set on the system, you will *not* be prompted to enter one. Use the `passwd` command to establish a password.

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

```
:/:<1>help
ftp      telnet
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
```

For more information on how to set up the `syslog` file and interpret it, see “Connecting to the Array” on page 100. For information on how to use the CLI commands, see the *Sun StorEdge T3+ Array Administrator’s Manual*.

▼ To Determine the Firmware Version

Perform this procedure if you are not sure which version of controller firmware your system is running. You determine the controller firmware version by logging onto the Sun StorEdge T3+ array through an Ethernet connection.

1. **Log on to the Sun StorEdge T3+ array controller firmware as described in “To Establish a Telnet Session” on page 101.**
2. **Type the following command:**

```
p2:/:<1> ver
```

An example output of the `ver` command for the master controller unit is shown in FIGURE 7-1. Check the firmware release version shown on the top line against the explanation shown in TABLE 7-1. In this case “2.00.01” represents firmware version 2.0.1 and an upgrade to version 2.1 is appropriate.

```
T3B Release 2.00.01 2001/01/04 15:18:16 (10.4.32.211)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.
```

FIGURE 7-1 Level 3 Firmware Version Check Output for Master Unit

TABLE 7-1 Explanation of `ver` Command Results

| <code>ver</code> command output: | |
|----------------------------------|--|
| T3B Release <i>x.xx.xx</i> | Equivalent Controller Firmware Version |
| 2.00 | 2.0 |
| 2.00.01 | 2.0.1 |
| 2.01.00 | 2.1—no upgrade required |

Upgrading the System Firmware

The controller firmware cannot be upgraded on an operational system. And for the upgrade to take effect, the controller must be reset (rebooted). While the controller reboots, the array is not available for storage.

This procedure must be done through an Ethernet connection.

The following conditions apply to this patch and firmware upgrade:

- The patch and firmware have to be resident on the host. The software is located on the SunSolve web site:

`http://sunsolve.sun.com`

- The Sun StorEdge T3+ array must have a root password prior to attempting this procedure.

This section contains the following topics:

- “To Install the System Firmware 2.1 Patch” on page 104
- “To Install the Interconnect Card Firmware” on page 107

▼ To Install the System Firmware 2.1 Patch

This procedure uses an uploading utility (script `t3.sh`) that automatically transfers the contents of this patch to the Sun StorEdge T3+ array. The script is only for use on a Solaris host system.

Caution – Before attempting to load firmware on a Sun StorEdge T3+ array, be sure to stop all I/O activity from all attached host systems.

Note – It is recommended that all Sun StorEdge T3+ array volumes be unmounted on Solaris before proceeding with this patch installation.

1. Download the system firmware 2.1 patch from the SunSolve site to the Solaris host.

Once the patch has been downloaded, transfer the contents to a temporary working directory if necessary.

2. Verify communication with the Sun StorEdge T3+ array to be upgraded.

```
$ ping IP_address
t3 is alive
```

Where *IP_address* is the T3+ IP address

3. Establish a telnet session to the Sun StorEdge T3+ array controller as described in “To Establish a Telnet Session” on page 101.

Note – The space available in the array’s reserved system area is limited. Therefore, it is important to ensure there is adequate space on the array before proceeding. It is not necessary to keep old images of controller firmware, unit interconnect card firmware, or EPROM binaries on the array after the upgrade.

4. Access the Solaris host directory that contains the firmware 2.1 patch.

5. Use the utility included in the patch (*t3.sh*) to download the files to the array by typing:

```
./t3.sh
```

Note – The *t3.sh* utility installs all controller and system files but does not overwrite the */etc/hosts* and */etc/syslog.conf* files on the target array system as these files are typically customized per local operating environment requirements. One exception is the battery refresh scheduler file */etc/schd.conf*. The *t3.sh* script makes a backup copy of this file on the array */etc/sch_old.conf* file. However, if required, factory default versions of these files are supplied in the *files.tar* archive image included in the patch.

The `t3.sh` utility prompts for the information shown in TABLE 7-2.

TABLE 7-2 `t3.sh` Utility Installation Prompts

| Prompt | Information Required |
|---|--|
| Please Enter Hostname or IP Address of Target T3+ To Be Upgraded: | Hostname or IP address of the target Sun StorEdge T3+ arra system |
| Please Enter Patch Location Pathname [.]: | Give the path to where the <code>files.tar</code> image exists. Typically the default current working directory is sufficient. Press ENTER to accept the current working directory. |
| Please Enter Your Home Directory Path [/home/joe_user]: | This is used by <code>ftp</code> to automatically load the patch file contents to the array. Give the home directory of the user ID used when launching the <code>t3.sh</code> script. |

6. Use the telnet session to verify that the files were transferred.

Two examples of files that should appear on the Sun StorEdge T3+ array system are:

```
:/:<1> cd /
:/:<2> ls
nb210.bin
lpc_05_02.bin
```

The `nb210.bin` and `lpc_05_02.bin` files are two examples of the expected files.

▼ To Install the Interconnect Card Firmware

The Interconnect card firmware is stored in the flash memory device on the interconnect card. The array can be operational during the interconnect card firmware upgrade.

1. Check the Interconnect Cards firmware version by typing:

```
:/: lpc version
      LOOP A      LOOP B
Enclosure 1  5.01 Flash  5.01 Flash
Enclosure 2  5.01 Flash  5.01 Flash
```

Contact your Sun service provider to find out the latest Interconnect card firmware version.

2. Use the fru stat command to find the Interconnect card port IDs.

These are listed under "LOOP."

```

:/: fru stat
CTLR      STATUS  STATE      ROLE      PARTNER    TEMP
-----  -
ulctr     ready   enabled    master    u2ctr      28.0
u2ctr     ready   enabled    alt master u1ctr      29.5

DISK      STATUS  STATE      ROLE      PORT1      PORT2      TEMP  VOLUME
-----  -
uld1      ready   enabled    data disk ready      ready      31    v1
uld2      ready   enabled    data disk ready      ready      31    v1
uld3      ready   enabled    data disk ready      ready      30    v1
uld4      ready   enabled    data disk ready      ready      31    v1
uld5      ready   enabled    data disk ready      ready      31    v1
uld6      ready   enabled    data disk ready      ready      31    v1
uld7      ready   enabled    data disk ready      ready      31    v1
uld8      ready   enabled    data disk ready      ready      31    v1
uld9      ready   enabled    data disk ready      ready      31    v1
u2d1      ready   enabled    data disk ready      ready      32    v2
u2d2      ready   enabled    data disk ready      ready      31    v2
u2d3      ready   enabled    data disk ready      ready      31    v2
u2d4      ready   enabled    data disk ready      ready      32    v2
u2d5      ready   enabled    data disk ready      ready      31    v2
u2d6      ready   enabled    data disk ready      ready      31    v2
u2d7      ready   enabled    data disk ready      ready      31    v2
u2d8      ready   enabled    data disk ready      ready      31    v2
u2d9      ready   enabled    data disk ready      ready      31    v2

LOOP      STATUS  STATE      MODE      CABLE      CABLE2     TEMP
-----  -
u211      ready   enabled    master    installed   -          26.0
u212      ready   enabled    slave     installed   -          28.0
u111      ready   enabled    master    -           installed  25.0
u112      ready   enabled    slave     -           installed  28.0

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

```

The Interconnect card port IDs are: u211, u212, u111, and u112.

3. Upgrade all Interconnect cards in an enterprise configuration by typing:

```
://: lpc download u211 lpc_05.02
://: lpc download u212 lpc_05.02
://: lpc download u111 lpc_05.02
://: lpc download u112 lpc_05.02
```

4. Verify the results by typing:

```
://: lpc version
                LOOP A           LOOP B
Enclosure 1    5.02 Flash        5.02 Flash
Enclosure 2    5.02 Flash        5.02 Flash
```

▼ To Upgrade the Controller Firmware

Caution – Before attempting to load firmware on a Sun StorEdge T3+ array, be sure to stop all I/O activity from all attached host systems.

This procedure upgrades both the master and alternate master controllers of a Sun StorEdge T3+ array enterprise configuration. This procedure also requires a Sun StorEdge T3+ array system reboot so that all necessary host preparations needed to sustain this procedure should be made before starting.

- 1. Establish a telnet session to the Sun StorEdge T3+ array controller as described in “To Establish a Telnet Session” on page 101.**
- 2. Set the boot mode to auto by typing:**

```
://: set bootmode auto
```

- 3. Initiate the download by typing:**

```
://: boot -i nb210.bin
```

Wait for the system to boot.

4. Reboot the Sun StorEdge T3+ array by typing:

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

Note – If, during the boot process, a controller detects a firmware version on the system disk different from the image loaded in flash, the controller will reflash its local level 3 image and reset. This can appear as two sequential boot cycles. This process is an expected behavior.

5. Repeat for each enterprise or workgroup configuration.

▼ To Verify the Upgrade

You can verify that the Sun StorEdge T3+ array upgrade is operating correctly by examining the output of several array commands.

Note – For more information on these commands, refer to the attached chapter, “Sun StorEdge T3+ Array Version 2.1 Controller Firmware CLI Commands.”

1. Type the following command:

```
p2:/:<1> ver
```

An example output of the `ver` command for the master controller unit is shown in FIGURE 7-2. Check the firmware release version shown on the top line against the explanation shown in TABLE 7-3. In this case, “2.01.00” represents firmware version 2.1 and indicates that the upgrade was successful.

```
T3B Release 2.01.00 2002/03/05 13:18:16 (10.4.32.211)
Copyright (C) 1997-2001 Sun Microsystems, Inc.
All Rights Reserved.
```

FIGURE 7-2 Level 3 Firmware Version Check Output for Master Unit

TABLE 7-3 Explanation of `ver` Command Results

| <code>ver</code> command output: T3B Release <i>x.xx.xx</i> | Equivalent Controller Firmware Version |
|--|--|
| 2.00 | 2.0 |
| 2.00.01 | 2.0.1 |
| 2.01.00 | 2.1—no upgrade required |

2. Use the `vol mode` command to verify that the volumes have been mounted, and that cache mirroring has been turned on.

The following is an example of the `vol mode` command output that shows volumes are mounted properly and cache mirroring is enabled in an enterprise configuration.

```
:/: vol mode

volume      mounted    cache      mirror
v0          yes        writebehind on
v1          yes        writebehind on
```

Note – For a workgroup configuration, cache mirroring and `mp_support` can be turned off.

3. Use the `port listmap` command to confirm that the volumes assigned to each array are correct.

```
:/: port listmap

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

4. Use the `fru stat` command to verify that the state of all components in the arrays is enabled.

The following is an example of `fru stat` command output from an enterprise configuration that indicates all components are operating properly:

```

:/: fru stat
CTLR      STATUS  STATE      ROLE      PARTNER    TEMP
-----
ulctr     ready   enabled    master    u2ctr      32.0
u2ctr     ready   enabled    alt master u1ctr      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
-----
ulpcu     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

```

5. Use the `fru list` command to make sure that all components in both arrays are detected by the host.

The following is an example of `fru list` command output from an enterprise configuration that shows all components detected by the host:

```

:/: fru list
ID      TYPE              VENDOR      MODEL      REVISION    SERIAL
-----
ulctr   controller card    0301        501-5710-01( 020100    101454
u2ctr   controller card    0301        501-5710-01( 020100    101125
uld1    disk drive         SEAGATE     ST318304FSUN A42D      3EL075C8
uld2    disk drive         SEAGATE     ST318304FSUN A42D      3EL06BYF
uld3    disk drive         SEAGATE     ST318304FSUN A42D      3EL083A9
uld4    disk drive         SEAGATE     ST318304FSUN 42D       3EL07MYA
uld5    disk drive         SEAGATE     ST318304FSUN A42D      3EL075PN
uld6    disk drive         SEAGATE     ST318304FSUN A42D      3EL073TH
uld7    disk drive         SEAGATE     ST318304FSUN A42D      3EL06YQ9
uld8    disk drive         SEAGATE     ST318304FSUN A42D      3EL07P62
uld9    disk drive         SEAGATE     ST318304FSUN A42D      3EL073F1
u2d1    disk drive         SEAGATE     ST318304FSUN A42D      3EL06TRE
u2d2    disk drive         SEAGATE     ST318304FSUN A42D      3EL07R81
u2d3    disk drive         SEAGATE     ST318304FSUN A42D      3EL06X3F
u2d4    disk drive         SEAGATE     ST318304FSUN A42D      3EL07X5M
u2d5    disk drive         SEAGATE     ST318304FSUN A42D      3EL06HX8
u2d6    disk drive         SEAGATE     ST318304FSUN A42D      3EL06096
u2d7    disk drive         SEAGATE     ST318304FSUN A42D      3EL07NWV
u2d8    disk drive         SEAGATE     ST318304FSUN A42D      3EL07PSX
u2d9    disk drive         SEAGATE     ST318304FSUN A42D      3EL07NXE
u1l1    loop card          SLR-MI      375-0085-01-5.02 Flash    028505
u1l2    loop card          SLR-MI      375-0085-01-5.02 Flash    028610
u2l1    loop card          SLR-MI      375-0085-01-5.02 Flash    029575
u2l2    loop card          SLR-MI      375-0085-01-5.02 Flash    030734
ulpcu1  power/cooling unitTECTROL-CAN 300-1454-01 ( 0000    026519
ulpcu2  power/cooling unitTECTROL-CAN 300-1454-01 ( 0000    026528
u2pcu1  power/cooling unitTECTROL 300-1453-01 ( 0001    01014
u2pcu2  power/cooling unitTECTROL-CAN 300-1454-01 ( 0000    039771
ulmpn   mid plane         SLR-MI      370-3990-01-0000    019191
u2mpn   mid plane         SLR-MI      370-3990-01-0000    022750

```

This completes the upgrade of the Sun StorEdge T3+ array controller firmware to version 2.1.

To enable volume slicing capability and to perform volume slicing, LUN masking, and mapping, see the *Sun StorEdge T3+ Array Administrator's Manual*.

Post-Installation Global Parameters Configurations

Before an array is shipped, the global parameters are set to default values. This chapter describes how to reconfigure your array by changing these default values.



Caution – If you are planning an enterprise configuration using new factory units, be sure to install and set up the units as a partner group before you power on, and change any parameters or create or change any logical volumes. Refer to Chapter 2 for more information.

Note – For more information on changing array global parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Topics in this chapter include:

- “Configuring Cache” on page 116
- “Configuring Logical Volumes” on page 119
- “Using RAID Levels to Configure Redundancy” on page 123

Configuring Cache

Each Sun StorEdge T3 array controller unit has a 256-Mbytes data cache; each Sun StorEdge T3+ array controller unit has 1 GByte of data cache. Writing to cache improves write performance by staging data in cache, assembling the data into data stripes, and then destaging the data from cache to disk, when appropriate. This method frees the data host for other operations while cache data is being destaged, and it eliminates the read-modify-write delays seen in noncache systems. *Read cache* improves performance by determining which data will be requested for the next read operation and prestaging this data into cache. RAID 5 performance is also improved by coalescing writes.

Performance and Redundancy

Cache mode can be set to the following values:

- **Auto.** The cache mode is determined as either write-behind or write-through, based on the I/O profile. If the array has full redundancy available, then caching operates in write-behind mode. If any array component is nonredundant, the caching mode is set to write-through. Read caching is always performed. Auto caching mode provides the best performance while retaining full redundancy protection.

Auto is the default cache mode for Sun StorEdge T3+ array.

- **Write-behind.** All write operations are cached. An algorithm determines when the data is destaged or moved from cache to disk. Write-behind cache improves performance, because a write to a high-speed cache is faster than a write to a normal disk.

Use write-behind cache mode with a workgroup configuration when you want to force write-behind caching to be used.



Caution – In a workgroup configuration, use a host-based mirroring solution to protect data. This configuration does not offer the redundancy to provide cache mirroring, and operating without a host-based mirroring solution could lead to data loss in the event of a controller failure.

- **Write-through.** This cache mode forces write-through caching to be used. In write-through cache mode, data is written through cache in a serial manner and is then written to the disk. Write-through caching does not improve write performance. However, if a subsequent read operation needs the same data, the read performance is improved, because the data is already in cache.

- **None.** No reads or writes are cached.

Note – For full redundancy in an enterprise configuration, set the cache mode and the mirror variable to Auto. This ensures that the cache is mirrored between controllers and that write-behind cache mode is in effect. If a failure occurs, the data is synchronized to disk, and then write-through mode takes effect. After the problem has been corrected and all internal components are again optimal, the system will revert to operating in write-behind cache mode.

Data Block Size Configuration

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 there are no volumes 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.

Note – The array data block size is independent of I/O block size. Alignment of the two is not required.

Data Block Size Selection

If the I/O initiated from the host is 4 Kbytes, a data block size of 64 Kbytes would force 8 Kbytes of internal disk I/O, wasting 4 Kbytes of the cache segment. Therefore, it would be best to configure 32-Kbyte block sizes, causing 4-Kbyte physical I/O from the disk. If sequential activity occurs, full block writes (32 Kbytes) will take place. For 8-Kbyte I/O or greater from the host, use 64-Kbyte blocks.

These applications benefit from the following data block or stripe unit sizes:

- 16-Kbyte data block size
 - Online Transaction Processing (OLTP)
 - Internet service provider (ISP)
 - Enterprise Resource Planning (ERP)
- 32-Kbyte data block size
 - NFS[™] file system, version 2

- Attribute-intensive NFS file system, version 3
- 64-Kbyte data block size
 - Data-intensive NFS file system, version 3
 - Decision Support Systems (DSS)
 - Data Warehouse (DW)
 - High Performance Computing (HPC)

Note – The data block size must be configured before any logical volumes are created on the units. Remember, this block size is used for every logical volume created on the unit. Therefore, it is important to have similar application data configured per unit.

Data block size is universal throughout a partner group. Therefore, you cannot change it after you have created a volume. To subsequently change the data block size, you must first delete the volume(s), change the data block size, and then create a new volume(s).



Caution – Unless you back up and restore the data on these volumes, it will be lost.

Mirrored Cache Enabling

By enabling mirrored cache, you can safeguard cached data if a controller fails.

Note – Mirrored cache is possible only in a redundant enterprise configuration.

Cache Allocation Configuration

Cache is allocated based on the read/write mix and it is dynamically adjusted by the controller firmware, based on the I/O profile of the application. If the application profile is configured for a 100 percent read environment, then 100% of the cache is used for reads. If the application profile has a high number of writes, then the upper limit for writes is set to 80 percent.

Configuring Logical Volumes

Also called a *logical unit number* (LUN), a logical volume is one or more disk drives that are grouped together to form a single unit. Each logical volume is represented to the host as a logical unit number. Using the `format` utility on the application host, you can view the logical volumes presented by the array. You can use this disk space as you would any physical disk—for example, to perform the following operations:

- Install a file system
- Use the device as a raw device (without any file system structure)
- Partition the device

Note – Individual physical disk drives are not visible from the application host. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more information on creating logical volumes.

Guidelines for Configuring Logical Volumes

Use the following guidelines when configuring logical volumes:

- The array's native volume management can support a maximum of two volumes per array unit.
- The minimum number of drives is based on the RAID level, as follows:
 - RAID 0 and RAID 1 require a minimum of two drives.
 - RAID 5 requires a minimum of three drives.
- Drive number 9 can be designated as a *hot spare*. If designated, drive number 9 will be the hot spare for all volumes in the array.
- A partial drive configuration is not allowed.
- Volumes cannot span array units.

Consider the following questions when configuring logical volumes:

- How many logical volumes do you need (one or two)?
- What RAID level do you require?
- Do you need a hot spare?

Determining the Number of Logical Volumes

You can configure a volume into seven partitions (also known as *slices*) using the `format(1M)` utility. Alternatively, you can configure virtually a large number of partitions (also known as *subdisks*) using VERITAS Volume Manager. Therefore, arrays are best configured as one large volume.

These applications benefit from the following logical volume or LUN configurations:

- Two LUNs per array
 - OLTP
 - ISP
 - ERP
 - NFS file system, version 2
 - Attribute-intensive NFS, version 3
- One LUN per array
 - Data-intensive NFS file system, version 3
 - DSS
 - DW
 - HPC

Note – If you are creating new volumes or changing the volume configuration, you must first manually rewrite the label of the previous volume using the `autoconfigure` option of the `format(1M)` UNIX host command. For more information on this procedure, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.



Caution – Removing and reconfiguring the volume will destroy all data previously stored there.

Determining the RAID Level

In general, RAID 5 is efficiently managed by the RAID controller hardware. This efficiency is apparent when compared to RAID 5 software solutions such as VERITAS Volume Manager.

The following applications benefit most from the RAID controller hardware of the array:

- Data-intensive NFS file system, version 3
- DSS
- DW
- HPC

Note – For more information about RAID levels, see “Using RAID Levels to Configure Redundancy” on page 123.

Determining a Hot-Spare

A hot-spare is always optional. Disk drive nine is always used as the hot-spare. If you choose to include a hot-spare in your configuration and you want it available for the first volume, you must specify it when you create the first volume in an array. This is also true if you create a second volume. If you want to add a hot spare at a later date to an existing volume, you must first delete that volume and then recreate it with the hot-spare designated during creation.

Note – Only one hot spare is allowed per array and it is only usable for the array in which it is configured. The hot spare must be configured as drive 9.

During a failure, for example, should drive 7 fail, drive 9 is synchronized automatically with the entire LUN to reflect the data on drive 7. After the failed drive (7) is replaced, the controller unit will automatically copy the data from drive 9 to the new drive, and drive 9 will become a hot spare again.

Tip – Although they are not required, hot spares are always recommended for mission-critical configurations because they allow the controller unit to reconstruct the data from the RAID group and only take a performance hit while the reconstruction is taking place. If a hot spare is not used, the controller unit remains in write-through cache mode until the failed drive is replaced and reconstruction is complete (which could take an extended period of time). During this time, the array is operating in degraded mode.

If no hot spare is present, the reconstruction of the data will begin when the failed drive is replaced—provided RAID 1 or RAID 5 is used.

Creating and Labeling a Logical Volume

You must set the RAID level and the hot-spare disk when you create a logical volume. For the Solaris operating system to recognize a volume, it must be labeled with the `format` or `fmthard` command.



Caution – Removing and reconfiguring a logical volume will destroy all data previously stored there.

Setting the LUN Reconstruction Rate

Note – When a failed drive is disabled, the volume is operating without further redundancy protection, so the failed drive needs to be replaced as soon as possible.

If the volume has a hot spare configured and that drive is available, the data on the disabled drive is reconstructed on the hot-spare drive. When this operation is complete, the volume is operating with full redundancy protection, so no loss of data occurs if another drive in the volume fails.

After a drive has been replaced, the original data is automatically reconstructed on the new drive. If no hot spare was used, the data is regenerated using the RAID redundancy data in the volume. If the failed drive data has been reconstructed onto a hot spare, after the reconstruction has completed, a copy-back operation begins in which the hot spare data is copied to the newly replaced drive.

You can also configure the rate at which data is reconstructed, so as not to interfere with application performance. Reconstruction rate values are low, medium, and high as follows:

- Low is the slowest and has the lowest impact on performance
- Medium is the default
- High is the fastest and has the highest impact on performance

Note – Reconstruction rates can be changed while a reconstruction operation is in process. However, the changes don't take effect until the current reconstruction has completed.

Using RAID Levels to Configure Redundancy

The RAID level determines how the controller reads and writes data and *parity* on the drives. The Sun StorEdge T3+ array can be configured with RAID level 0, RAID level 1 (1+0), or RAID level 5. The factory-configured LUN is a RAID 5 LUN.

Note – The default RAID level (5) can result in very large volumes—for example, 128 Gbytes in a configuration of a single 7+1 RAID 5 LUN plus a hot spare, with 18-Gbyte drives. Some applications cannot use such large volumes effectively. The following three solutions can be used separately or in combination:

- Use the partitioning utility available on the data host's operating system. In the Solaris environment, use the `format` utility, which can create up to seven distinct partitions per volume. Note that in the case of the configuration described above, if each partition is equal in size, this will result in 18-Gbyte partitions, which still may be too large to be used efficiently by legacy applications.
- Use third-party software on the host system to create as many partitions as desired from a given volume. In the Solaris environment, you can use VERITAS Volume Manager or Solaris Logical Volume Management (SLVM)—formerly known as Solstice DiskSuite (SDS)—for this purpose.
- Use volume slicing.

Note – For information on using the `format` utility, refer to the `format (1M)` man page. For more information on third-party software or VERITAS Volume Manager, refer to the documentation for that product.

RAID 0

Data blocks in a RAID 0 volume are striped across all the drives in the volume in order. There is no parity data, so RAID 0 uses the full capacity of the drives. There is, however, no redundancy. If a single drive fails, all data on the volume is lost.

TABLE 8-1 shows the data layout in a RAID 0 volume.

TABLE 8-1 RAID 0 Data Blocks¹

| Stripe | Drive1 | Drive2 | Drive3 | Drive4 | Drive5 | Drive6 | Drive7 | Drive8 | Drive9 |
|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| 0 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 1 | 9 | A | B | C | D | E | F | 10 | 11 |
| 2 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 1A |

1. Block numbers are written in hexadecimal notation.

RAID 1

Each data block in a RAID 1 (1+0) volume is mirrored on two drives, and the blocks are striped across all the drives in the volume. If one of the mirrored pair fails, the data from the other drive is used. Because the data is mirrored in a RAID 1 configuration, the volume has only half the capacity of the assigned drives. For example, if you create a four-drive RAID 1 volume with 18-Gbyte drives, the resulting data capacity is $2 \times 18 = 36$ Gbytes. TABLE 8-2 shows the data blocks in a RAID 1 (1+0) volume with nine drives.

TABLE 8-2 RAID 1 (1+0) Data Blocks¹

| Stripe | Drive1 | Drive2 | Drive3 | Drive4 | Drive5 | Drive6 | Drive7 | Drive8 | Drive9 |
|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| 0 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 1 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 0 |
| 2 | 9 | A | B | C | D | E | F | 10 | 11 |

1. Block numbers are written in hexadecimal notation.

RAID 5

In a RAID level 5, data is striped across the drives in the volumes in segments, with parity information striped across the drives, as well. Because of this parity, if a single drive fails, data can be recovered from the remaining drives. Two drive failures cause all data to be lost. A RAID 5 volume has the data capacity of all the drives in the LUN, less one. For example, a five-drive RAID 5 volume with 18-Gbyte drives has a data capacity of $4 \times 18 = 72$ Gbytes. TABLE 8-3 shows the data blocks in a RAID 5 volume. The location of the parity drive in the stripe rotates from right to left in order to eliminate a performance bottleneck.

TABLE 8-3 RAID 5 Data Blocks¹

| Stripe | Drive1 | Drive2 | Drive3 | Drive4 | Drive5 | Drive6 | Drive7 | Drive8 | Drive9 |
|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| 0 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | P |
| 1 | 9 | A | B | C | D | E | F | P | 8 |
| 2 | 12 | 13 | 14 | 15 | 16 | 17 | P | 10 | 11 |
| 3 | 1B | 1C | 1D | 1E | 1F | P | 18 | 19 | 20 |
| 4 | 24 | 25 | 26 | 27 | P | 20 | 21 | 22 | 23 |
| 5 | 2D | 2E | 2F | P | 28 | 29 | 2A | 2B | 2C |
| 6 | 36 | 37 | P | 30 | 31 | 32 | 33 | 34 | 35 |
| 7 | 3F | P | 38 | 39 | 3A | 3B | 3C | 3D | 3E |
| 8 | P | 40 | 41 | 42 | 43 | 44 | 45 | 46 | 47 |
| 9 | 48 | 49 | 4A | 4B | 4C | 4D | 4E | 4F | P |

1. Block numbers are written in hexadecimal notation.

Configuring RAID Levels

The Sun StorEdge T3+ array is preconfigured at the factory with a single LUN, RAID level 5 redundancy, and no hot spare. After a volume has been configured, you cannot reconfigure it to change its size, RAID level, or hot-spare configuration. You must first delete the volume and create a new one with the configuration values you want.

Post-Installation Array Settings Configuration

This chapter describes how to reconfigure the default settings of the array as well as how to monitor the array.

Topics in this chapter include:

- “Reconfiguring the Array Settings” on page 128
- “Monitoring the Array” on page 132

Reconfiguring the Array Settings

The array is set with the following default configuration:

- A RAID volume configuration
- Cache mode set to auto
- Disk drives set to autodisable and autoreconstruct

These default settings and other possible configurations are described in the following sections. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for instructions on how to change the array default settings.

Some of the terminology used in this chapter is defined as follows:

- *Volume*, also known as a *logical unit number* (LUN), is the basic structure you create across the drive group to retrieve and restore data. A volume is a set of physical drives that has been configured into RAID level 0, 1, or 5.
- A *hot-spare* is a drive containing no data that acts as a standby in case another drive fails in a RAID 1 or 5 volume. The hot-spare drive adds another level of redundancy: If a drive fails, the hot-spare takes over for the failed drive until you replace it. When you replace the failed drive, data from the hot-spare is copied back to the replaced drive and the hot-spare returns to a standby status.
- *Parity* is additional information stored along with the data that enables the controller to rebuild data after a drive failure. Only RAID 5 volumes have parity. RAID 1 uses mirroring, not parity, which enables the controller to reconstruct data after a drive failure.

Unit Volume Configurations

The drives in an array unit can be grouped into one or two volumes (LUNs). The minimum number of drives to compose a volume is two; the maximum is nine. A hot-spare drive can be configured in the unit as an option. If you are configuring a hot-spare in the volume, you can use only drive 9 in the array as a hot-spare. You can have only one hot-spare per array, but it can be shared by two volumes.

If you choose to include a hot-spare disk drive in your configuration, add it when you initially create the volume. If you want to add a hot-spare at a later date, you must remove the existing volume and re-create the configuration.

You can also configure the two volumes into smaller partitions or combine multiple volumes into a single volume by using specialized third-party software such as VERITAS Volume Manager.

Note – If you are creating new volumes or changing the volume configuration, you must first manually rewrite the label of the previous volume by using the `autoconfigure` option of the `format(1M)` utility. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more information on this procedure.

Examples of valid volume configurations are shown in TABLE 9-1.

TABLE 9-1 Examples of Volume Configurations

| 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 | |

RAID Levels

See “Using RAID Levels to Configure Redundancy” on page 123.

Cache Modes

Data cache is used for the following purposes:

- **Buffering.** Data that is being transferred between the host and the drives.
- **Read caching.** Data for future retrieval, to reduce disk I/O as much as possible.
- **Write caching.** Data that is used to build up stripes of data, eliminating the read-modify-write overhead. Write caching improves performance for applications that are writing to disk.

Four cache modes are supported: auto, write-behind, write-through, and none. The default mode for the Sun StorEdge T3+ Array is auto.

- In the *auto* cache mode, the caching mode is write-behind if the array has full redundancy available. If any component is nonredundant, the cache mode is set to write-through. Read caching is always performed. This mode provides the best performance while retaining full redundancy protection.
- The *write-behind* cache mode forces write-behind caching to always be used. Data loss may occur in this mode if the units are not configured as a fully redundant enterprise and an array system failure occurs.



Caution – To avoid data loss in the write-behind cache mode, make sure you use host-based mirroring.

- The *write-through* cache mode forces write-through caching to always be used. Read caching is also performed in this mode.
- The *none* cache mode forces no write caching and no read caching.



Caution – If one of the redundant components fail, the cache mode is set to write-through. In the write-through mode, mirroring is not used and data is not being cached. If you were to view the cache mode at this time using the `vol mode` command, the setting for `cache` would be displayed as `writethrough` and the setting for `mirror` would be displayed as `on`. This means that mirroring is enabled, but it is *not* being used.

The following table highlights the system states that can disable write-behind cache. In cases where write-behind cache has been disabled, the array is designed to protect the data in cache. If any hardware component fails that might compromise the safety of this data, the array will disable write-behind cache, causing outstanding cached data to be flushed to disk. Write-behind caching is re-enabled when the failed condition has been corrected.

TABLE 9-2 Cache Transition States

| Failure or Problem | Write-Behind Cache State |
|----------------------------|--------------------------|
| Disk hard failure (RAID 5) | Disabled |
| Disk hard failure (RAID 1) | Disabled |
| Disk hard failure (RAID 0) | Disabled |
| Disk reconstruct | Disabled |
| Disk copy-back operation | Disabled |
| Interconnect card 1 | Disabled |

TABLE 9-2 Cache Transition States *(Continued)*

| Failure or Problem | Write-Behind Cache State |
|----------------------------|---|
| Interconnect card 2 | Disabled |
| Interconnect card 1+2 | Disable (system down) |
| Controller | Disable (controller failover) |
| Battery | Disabled |
| Battery warranty expires | Disabled |
| Battery refresh | Disable (hold time low) |
| 28-day battery refresh | Disable (during 12-minute battery discharge) ¹ |
| Power and cooling unit 1 | Disabled |
| Power and cooling unit 2 | Disabled |
| Power and cooling unit 1+2 | Disable (system down) |
| Battery DC not OK | Disabled |
| Host cable | Enable (LUN failover) |
| Media Interface Adapter | Enable (LUN failover) |
| Host GBIC | Enable (LUN failover) |
| Missing FRU | Disabled |

1. Battery discharge may be slightly longer or shorter than 12 minutes.

Disabling and Reconstructing a Drive

If a drive fails, it is automatically disabled by the system. This prevents any further access to the drive. The data on the drive is available by reconstructing it using RAID redundancy data in the volume. When a drive is disabled, the volume is operating without any further redundancy protection, so the failed drive needs to be replaced as soon as possible.

If the volume has a hot-spare configured and that drive is available, the data on the disabled drive is reconstructed on the hot-spare drive. When this operation is complete, the volume is operating with full redundancy protection, so another drive in the volume may fail without loss of data.

After a drive has been replaced, the original data is automatically reconstructed on the new drive. If no hot-spare was used, the data is regenerated using the RAID redundancy data in the volume. If the failed drive data has been reconstructed onto a hot-spare, After the reconstruction has completed, a copy-back operation begins during which the hot-spare data is copied to the newly replaced drive.

Monitoring the Array

You can monitor the array remotely by using any of the following tools:

- Simple Network Management Protocol (SNMP) notification
- `syslog` error reports

SNMP Notification

SNMP notification directs the software to send remote notification of array events to the administrator's host by using SNMP traps. An SNMP trap is a UDP datagram sent to a set of hosts that are identified by their IP addresses. The contents of the trap message include header information that identifies the type of trap and the data associated with the trap message.

When an event occurs in the array unit that requires attention, a trap is asynchronously sent to the host to report the event. The host(s) to which a trap is sent is configured by the user. The host software that receives the trap is an SNMP management package.

Sun StorEdge T3+ arrays use the `/etc/syslog.conf` file to trigger the SNMP trap. This file contains four levels of error warning to the administrator. For instructions on how to set up SNMP notification for the array, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

An SNMP MIB file is provided with the array for use with SNMP-based applications. The MIB file is located in the following directory on the array:

```
/web/snmp/t300.mib
```

To use this file, use ftp to transfer this MIB file from the array to your SNMP management host and refer to your specific SNMP management documentation for correct MIB file installation procedures.

syslog Error Reports

A `syslog` daemon exists in the hardware RAID controller that records system messages and provides for remote monitoring. There are four possible levels of messages.

TABLE 9-3 `syslog` Error Reporting Levels

| Message Level | Indication | Example |
|--------------------|--|--|
| Error | Critical system event requiring immediate user intervention or attention. | An over-temperature condition or a removed FRU |
| Warning | A system event that will eventually require user intervention. | A FRU being disabled and a recovery procedure being executed |
| Notice | A system event that might be a side effect of other events or possibly a normal condition. | The power switch turned off |
| Information | A system event that has no effect on the health of the system. | The healthy state of a FRU |

You can configure the level of error reporting to receive only certain types of messages, such as error and warning messages. For more information on configuring remote system logging, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Post-Installation Partner Groups Configuration

This chapter describes how to configure array partner groups. Sun StorEdge T3+ arrays can be interconnected in partner groups to form a redundant and larger storage system.

Note – The terms *partner group* and *enterprise configuration* refer to the same type of configuration and are used interchangeably in this document.

Note – Partner groups are *not* supported in Sun Cluster 2.2 configurations.

Topics in this chapter include:

- “Understanding Partner Groups” on page 136
- “How Partner Groups Work” on page 137
- “Creating Partner Groups” on page 138

Understanding Partner Groups

In a partner group, there is a *master controller unit* and an *alternate master controller unit*. The master controller unit is the array positioned at the bottom of an array stack in either a rackmounted or tabletop installation. The alternate master controller unit is positioned on top of the master controller unit. Array units are connected using the interconnect cards and interconnect cables. A partner group is shown in FIGURE 10-1.




FIGURE 10-1 Sun StorEdge T3+ Array Partner Group

Note – The Sun StorEdge T3+ array configurations support direct FC-AL connections.

When two units are connected together, they form a redundant partner group. This group provides controller redundancy. Because the controller is a single point of failure in a standalone configuration, this redundancy allows an application host to access data even if a controller fails. This configuration offers multipath and LUN failover features.

The partner group connection also allows for a single point of control. The bottom unit will assume the role of the master, and from its Ethernet connections, it will be used to monitor and administer the unit installed above it.

The master controller unit will set the global variables within this storage system, including cache block size, cache mode, and cache mirroring.

Note – For information about setting or changing these parameters, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Any controller unit will boot from the master controller unit's drives. All configuration data, including `syslog` information, is located on the master controller unit's drives.

How Partner Groups Work

If the master controller unit fails and the “heartbeat” between it and the alternate master stops, this failure causes a controller failover, so the alternate master assumes the role of the master controller unit. The new master (formerly the alternate master) takes the IP address and the MAC address from the old master and begins to function as the administrator of the storage system. It will also be able to access the drives of the former master controller unit. Those drives will still be used to store `syslog` information, system configuration information, and bootcode. Should it become necessary to reboot the storage system while the master controller unit is inactive, the alternate master will use the former master controller unit's drives to boot.

Note – Even after the failed master controller is back online, it remains the alternate master controller and, as a result, the original configuration is modified from its original state.

In a redundant partner group configuration, the units can be set to do a path failover operation. Normally the volumes or LUNs that are controlled by one unit are not accessible to the controller of the other. The units can be set so that if a failure in one controller occurs, the remaining one will accept I/O for the devices that were running on the failed controller. To enable this controller failover operation, multipathing software—such as VERITAS Volume Manager, Sun StorEdge Traffic Manager software, or Solaris Alternate Pathing (AP) software—must be installed on the data application host.

Note – In order for a feature such VERITAS DMP to access a LUN through both controllers in a redundant partner group, the `mp_support` parameter must be set to `rw` to enable this feature. If you are using Sun StorEdge Traffic Manager, the `mp_support` parameter must be set to `mpxio`. For information on setting the `mp_support` parameter and options, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Creating Partner Groups

Partner groups can be created in two ways:

- From new units
- From existing standalone units

Instructions for installing new array units and connecting them to create partner groups can be found in Chapter 2.

To configure existing standalone arrays with data into a partner group, you must go through a qualified service provider. Contact your Sun Service representative for more information.



Caution – The procedure to reconfigure the arrays into a partner group involves deleting all data from array disks and restoring the data after the completing the reconfiguration. There is the potential risk of data loss or data corruption if the procedure is not performed properly.

Post-Installation Host Connections

This chapter describes the host bus adapters (HBAs) that are supported by Sun StorEdge T3+ arrays.

Topics in this chapter include:

- “Sun Enterprise SBus+ and Graphics+ I/O Boards” on page 140
- “Sun StorEdge PCI FC-100 Host Bus Adapter” on page 141
- “Sun StorEdge SBus FC-100 Host Bus Adapter” on page 142
- “Sun StorEdge PCI Single Fibre Channel Network Adapter” on page 143
- “Sun StorEdge PCI Dual Fibre Channel Network Adapter” on page 144
- “Compact PCISun StorEdge CompactPCI Dual Fibre Channel Network Adapter” on page 145

Sun Enterprise SBus+ and Graphics+ I/O Boards

The SBus+ and Graphics+ I/O boards each provide mounting for two *Gigabit Interface Converters* (GBICs). For more detailed information about these I/O boards, refer to the *Sun Enterprise 6x00/5x00/4x00/3x00 Systems SBus+ and Graphics+ I/O Boards Installation Guide*. FIGURE 11-1 shows a Sun Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O board.




FIGURE 11-1 Sun Enterprise 6x00/5x00/4x00/3x00 SBus+ I/O Board

System Requirements

Your system must meet the following hardware and software requirements:

- Sun Enterprise 6x00/5x00/4x00/3x00 system
- An available I/O board slot
- OpenBoot™ PROM, version 3.2.10 or later
- A release of the Solaris operating environment that supports this board. The first release that supports this board is the Solaris 2.6 operating environment.

Sun StorEdge PCI FC-100 Host Bus Adapter

The Sun StorEdge PCI FC-100 host bus adapter is a 33-MHz, 100 Mbytes/second, single-loop Fibre Channel PCI host bus adapter with an onboard GBIC. This host bus adapter is PCI Version 2.1-compliant. For more detailed information about this product, refer to the *Sun StorEdge PCI FC-100 Host Adapter Installation Manual*. FIGURE 11-2 shows a Sun StorEdge PCI FC-100 host bus adapter.




FIGURE 11-2 Sun StorEdge PCI FC-100 Host Bus Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available PCI port
- A release of the Solaris operating environment that supports this board. The first release that supports this board is the Solaris 2.6 operating environment

Sun StorEdge SBus FC-100 Host Bus Adapter

The Sun StorEdge SBus FC-100 host bus adapter is a single-width Fibre Channel SBus card with a Sun Serial Optical Channel (SOC+) application-specific integrated circuit (ASIC). You can connect up to two loops to each card, using hot-pluggable GBICs. For more detailed information about this product, refer to the *Sun StorEdge SBus FC-100 Host Adapter Installation and Service Manual*. FIGURE 11-3 shows a Sun StorEdge SBus FC-100 host bus adapter.




FIGURE 11-3 Sun StorEdge SBus FC-100 Host Bus Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available SBus port
- A release of the Solaris operating environment that supports this board. The first release that supports this board is the Solaris 2.6 operating environment.

Sun StorEdge PCI Single Fibre Channel Network Adapter

The Sun StorEdge PCI Single Fibre Channel network adapter is a Fibre Channel PCI card with one onboard optical receiver. This network adapter is PCI Version 2.1-compliant. For more detailed information about this product, refer to the *Sun StorEdge PCI Single Fibre Channel Network Adapter Installation Guide*. FIGURE 11-4 shows a Sun StorEdge PCI Single Fibre Channel network adapter.




FIGURE 11-4 Sun StorEdge PCI Single Fibre Channel Network Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available PCI port
- A release of the Solaris operating environment that supports this board. The first release that supports this board is the Solaris 7 11/99 operating environment.

Sun StorEdge PCI Dual Fibre Channel Network Adapter

The Sun StorEdge PCI Dual Fibre Channel network adapter is a Fibre Channel PCI card with two onboard optical transceivers. This network adapter is PCI Version 2.1-compliant. For more detailed information about this product, refer to the *Sun StorEdge PCI Dual Fibre Channel Network Adapter Installation Guide*. FIGURE 11-6 shows a Sun StorEdge PCI Dual Fibre Channel network adapter.




FIGURE 11-5 Sun StorEdge PCI Dual Fibre Channel Network Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available PCI slot
- A release of the Solaris operating environment that supports this board. The first release that supports this board is the Solaris 7 11/99 operating environment.

Compact PCISun StorEdge CompactPCI Dual Fibre Channel Network Adapter

The Sun StorEdge CompactPCI Dual Fibre Channel network adapter has two 1-Gbit Fibre Channel ports on a cPCI card. For more detailed information about this product, refer to the *Sun StorEdge CompactPCI Dual Fibre Channel Network Adapter Installation Guide*. FIGURE 11-6 shows a Sun StorEdge CompactPCI Dual Fibre Channel network adapter.




FIGURE 11-6 Sun StorEdge CompactPCI Dual Fibre Channel Network Adapter

System Requirements

Your system must meet the following hardware and software requirements:

- An available cPCI port
- OpenBoot PROM version 5.1 or later
- Solaris 8 operating environment

Post-Installation Array Cabling

This chapter describes the array configurations supported by the Sun StorEdge T3+ arrays.

Topics in this chapter include:

- “Overview of Array Cabling” on page 148
- “Workgroup Configurations” on page 150
- “Enterprise Configurations” on page 151

Overview of Array Cabling

Sun StorEdge T3+ arrays supports the following connections:

- One FC-AL interface to the application host
- One Ethernet interface to the management host (by way of a LAN) for administration purposes
- One serial interface to be used for service tasks by qualified service personnel only
- Interconnect ports for configuring arrays into partner groups

Data Path

For the data path (FC-AL) connection, there are three ways that the array can connect to the host:

- Direct attached mode to the data host
- Hub connection, where the FC-AL from the array is connected to a hub on the same network as the data host
- Switch connection, where the FC-AL from the array is connected to a switch on the same network as the data host

Administration Path

For the administration path, each controller unit has an Ethernet connector. For each installed controller, an Ethernet connection and IP address are required.

Note – In a partner group, only one of the two Ethernet connections is active at any time. The second Ethernet connection is used for redundancy.

Connecting Arrays

The array also has two interconnect cards that connect the array in an enterprise configuration. These interconnect cards have two ports (in and out).

Note – An array enterprise configuration is *not* supported by Sun Cluster 2.2.

Note – In a workgroup (standalone) configuration, these interconnect cards cannot be used to connect to the administrative console or to the application host. These interconnect cards are used solely for ensuring redundancy and failover mechanisms in partner groups.

FIGURE 12-1 shows a Sun StorEdge T3+ array with a controller card and interconnect cards.




FIGURE 12-1 Sun StorEdge T3+ Array Controller Card and Interconnect Cards

Workgroup Configurations

The following configuration rules apply to array workgroup configurations (FIGURE 12-2):

- The interconnect ports, which are used only in partner group configurations, cannot be used for host connections.
- The FC-AL connection provides a data path to the application host.
- The Ethernet connection provides a link to the management host.
- The serial port is used solely for diagnostics and service by qualified service personnel only.
- Fiber-optic cable lengths between 2 and 500 meters, using short-wave laser and 50-micron fiber-optic cable, are supported.




FIGURE 12-2 Array Workgroup Configuration

Enterprise Configurations

The following rules configuration rules apply to enterprise (partner group) configurations (FIGURE 12-3):

- The interconnect ports, which are used only in enterprise configurations, cannot be used for host connections.
- The FC-AL connection provides a data path to the application host.
- The Ethernet connection provides a link to the management host.
- The serial port is used solely for diagnostics and service by qualified service personnel only.
- Fiber-optic cable lengths between 2 and 500 meters, using short-wave laser and 50-micron fiber-optic cable, are supported.

This configuration is optimal because it provides full redundancy to the application hosts. Failover mechanisms are provided within the arrays, but the application host has to provide data-path failover mechanisms, such as Dynamic Multi-Pathing from VERITAS Volume Manager or Alternate Pathing from Sun Enterprise Server Alternate Pathing.




FIGURE 12-3 Enterprise Configuration

Troubleshooting the Array

This chapter describes how to approach troubleshooting the array. Use the information in this chapter along with the *Sun StorEdge T3+ Array Administrator's Manual* to determine the service required.

Topics in this chapter include:

- “ofdg Notice” on page 154
- “Analyzing the Problem” on page 155
- “Channel Connection Failures” on page 161
- “FRU Failures” on page 161

ofdg Notice



Note – Sun StorEdge T3+ array `ofdg` utility is for service use only. This utility makes data unavailable.

The Sun StorEdge T3+ array firmware provides a command-line and graphical user interface. The `ofdg` utility is used to perform offline diagnostic testing. This tool should be used *only* by qualified Sun service providers at this time. The `ofdg` utility is limited in its test capabilities and must be run while the system is offline to avoid problems.

Analyzing the Problem

To begin the problem analysis, check one or a combination of the following four resources, as described in the following sections.

- “Host-Generated Message” on page 155
- “Telnet Session” on page 155
- “Array LEDs” on page 156

Host-Generated Message

The data host sees the array as one or two logical volumes. A host-generated message might indicate that the host cannot communicate with the array through the FC-AL channel, or that an excessive number of channel errors are occurring. If the host loses access to the array through the channel connection, then any host messages regarding the array will refer only to the logical volume(s). To see details of problems internal to the array, use the tool(s) that you installed on your management console.

In an enterprise configuration where Alternate Pathing failover has been established, the failure of a channel path or array controller causes the host to redirect I/O from the failed channel to the second FC-AL connection.

Telnet Session

The Telnet session is a direct network link to the array unit by means of the command-line interface (CLI). You can run individual commands to query and repair the unit from this interface. The Telnet session requires access to the unit’s IP address. You can run the CLI from any host that is part of the same network.

You can quickly determine the status of the array from the CLI. You can also examine the `syslog` file of the array file system through the CLI. The `syslog` file contains a record of events that have occurred in the unit.

For more information on how to set up the `syslog` file and interpret it and for information on how to use the CLI commands, see the *Sun StorEdge T3+ Array Administrator’s Manual*.

Array LEDs

Each FRU component within the array (except for the chassis) has a set of LEDs that indicate health and status, as shown in the following sections:

- “Disk Drive LEDs” on page 156
- “Power and Cooling Unit LEDs” on page 157
- “Interconnect Card LEDs” on page 159
- “Controller Card LEDs” on page 160

Even though the LED indicates a FRU failure, you should always verify the FRU status using either the CLI before replacing the component. Refer to the *Sun StorEdge T3+ Array Administrator’s Manual* for instructions.

Note – See Chapter 14 for instructions on how to service a failed FRU.

Disk Drive LEDs

The top of each disk drive has LEDs that indicate drive activity and status. These LEDs can be seen with the front cover on the unit. TABLE 13-1 lists the possible drive LED states and describes each state. See FIGURE 13-1.




FIGURE 13-1 Disk Drive LEDs (Viewed Through Front Cover)

TABLE 13-1 Disk Drive LED Descriptions

| Drive Activity (green) | Drive Status (amber) | Description |
|------------------------|----------------------|---|
| Off | Off | Drive is not installed (not recognized) |
| Slow blink | Off | Drive is spinning up or down |
| Solid | Off | Drive is OK and is idle |
| Flashing | Off | Drive is OK and is active |
| n/a | Solid | A drive reconstruct or firmware download is in progress |
| n/a | Slow blink | Drive failure has occurred; it is OK to replace drive |

Note – See Chapter 14 for instructions on how to service a failed disk drive.

Power and Cooling Unit LEDs

Each power and cooling unit has an AC LED and a power supply (PS) LED. TABLE 13-2 lists the possible conditions of these LEDs and describes each state. See FIGURE 13-2.




FIGURE 13-2 Power and Cooling Unit LEDs

TABLE 13-2 Power and Cooling Unit LED Descriptions

| AC LED (green or amber) | PS LED (green or amber) | Description |
|-------------------------|-------------------------|---|
| Off | Off | Power is off No AC power is available |
| Amber | Off | Power is off Power switch turned off AC power is available |
| Green | Off | Occurs when array is shut down: <ul style="list-style-type: none"> • Power and cooling unit is disabled • AC power is available |
| Green | Green | Normal operating state: <ul style="list-style-type: none"> • Power and cooling unit is receiving AC power • Power switch is turned on • AC power is available |
| Amber | Amber | Switch is off; array will power off after power and cooling unit is disabled |
| Green | Amber | Indicates one or more of the following: <ul style="list-style-type: none"> • Over-temperature condition; PCU disabled • DC power not available; PCU disabled • Both fans fault; PCU disabled • Battery on refresh cycle |
| Green | Blinking green | Battery on refresh cycle |
| Green | Blinking amber | Indicates one or more of the following: <ul style="list-style-type: none"> • PCU disabled • One fan fault • Battery hold-time low; PCU remains enabled • Battery out of warranty; PCU remains enabled • Battery life-span failure; PCU remains enabled |

Note – See Chapter 14 for instructions on how to service a failed power and cooling unit.

Interconnect Card LEDs

Each interconnect card has a status LED for each interconnect cable. TABLE 13-3 lists the possible interconnect card status LED states with descriptions of each state. See FIGURE 13-3.




FIGURE 13-3 Interconnect Card LEDs

TABLE 13-3 Interconnect Card LED Descriptions

| Interconnect Card Status LED (green or amber) | Description |
|---|--|
| Off | Interconnect card not installed (not recognized) |
| Green-solid | Interconnect card OK Cable OK (if present) |
| Green-slow blink | Interconnect card OK Cable bad; OK to replace cable |
| Amber-solid | Interconnect card firmware download in progress |
| Amber-slow blink | Interconnect card failure; OK to replace interconnect card |

Note – See Chapter 14 for instructions on how to service a failed interconnect card.

Controller Card LEDs

The Sun StorEdge T3+ array controller card LEDs are described in the following section.

Sun StorEdge T3+ Array Controller Card LEDs

The Sun StorEdge T3+ array controller card has two channel-active LEDs—one for the FC-AL interface port and one for the Ethernet port—and a controller status (online) LED. TABLE 13-4 lists the possible states of the controller card LEDs and describes each state.




FIGURE 13-4 Sun StorEdge T3+ Array Controller Card LEDs

TABLE 13-4 Sun StorEdge T3+ Array Controller Card LED Descriptions

| LED | Action | Description |
|---|----------------|---------------------------------------|
| FC-AL Channel Active LED (green) | Off | The port is disabled |
| | Green | The port is enabled and idle |
| | Blinking green | The port is enabled and active |
| Ethernet Active LED (green) | Off | Link invalid |
| | Green | Link valid and idle |
| | Blinking green | Link valid and active |
| 10/100 BASE-T Active LED (green) | Off | Port disabled (10 Mbps rate) |
| | Green | Port enabled and idle (100 Mbps rate) |
| | Blinking green | Port enabled and active |

TABLE 13-4 Sun StorEdge T3+ Array Controller Card LED Descriptions

| LED | Action | Description |
|---|----------------|---|
| Controller Status LED (green or amber) | Off | Controller not installed (not recognized) |
| | Green | Controller OK |
| | Amber | Controller boot, shutdown, or firmware download in progress |
| | Blinking amber | Controller failure; OK to replace controller |

Channel Connection Failures

Host-to-array channel failures occur when the connection between the array and the host is either severed or intermittent. The components that make up this data channel connection are:

- Host bus adapter (HBA), which resides on the host
- GBIC adapter, used to connect the FC-AL cable to an SBus HBA
- Fibre Channel cable that connects the array to the host
- Channel interface port in the array

Errors in the host data channel are outside the scope of the 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.

FRU Failures

FRU failures are reported by an alarm or message notification from the tool used to monitor the array. The failed FRU will most likely display a blinking amber LED. Before replacing a FRU based on the LED indicator, verify its failure by using the command-line interface (CLI).

To check for a FRU failure using the CLI, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Servicing the Array

This chapter contains information about removing and replacing the field-replaceable units (FRUs) in an array and upgrading software and firmware.



Caution – Any FRU that is removed must be replaced within 30 minutes or the array and all attached arrays will automatically shut down and power off.

Topics in this chapter include:

- “ofdg Notice” on page 164
- “Preparing for Service” on page 165
- “Removing and Replacing Components” on page 165

ofdg Notice



Note – Sun StorEdge T3+ array `ofdg` utility is for service use only. Use of this feature will make data unavailable.

The Sun StorEdge T3+ array firmware provides a command-line and graphical user interface to a diagnostics tab provided by Sun StorEdge Component Manager 2.2. The `ofdg` utility is used to perform offline diagnostic testing. This tool should be used *only* by qualified Sun service providers at this time. The `ofdg` utility is limited in its test capabilities and must be run while the system is offline to avoid problems.

Preparing for Service

Follow these procedures to help prevent damaging FRUs:

- Wear a grounding wrist strap.
- Before handling a FRU, discharge any static electric charge by touching a ground surface.
- Do not remove a FRU from its antistatic protective bag until it is ready for installation.
- When removing a FRU from the array, immediately place it in an antistatic bag and packaging.
- Handle a FRU only by its edges and avoid touching the circuitry.
- Do not slide a FRU over any surface.
- Avoid having plastic, vinyl, and foam in the work area.
- Limit body movement (which builds up static electricity) during FRU installation.



Caution – Although the array will remain operational for up to 30 minutes with a FRU removed, replace a failed FRU as soon as you have removed it from the array.

Removing and Replacing Components

This section contains instructions for servicing the following array FRUs:

- “Disk Drives” on page 166
- “Power and Cooling Units” on page 170
- “UPS Battery” on page 173
- “Interconnect Cards” on page 174
- “Controller Card” on page 177
- “Chassis” on page 179

Disk Drives



Caution – Replace only one disk drive in an array at a time to ensure that no data is lost. Ensure that the disk drive is fully reenabled before replacing another disk drive in the same array.

The default configuration for the array is to automatically spin up and reenble a replaced disk drive, then automatically reconstruct the data from the parity or hot-spare disk drives. Disk drive spinup takes about 30 seconds, and reconstruction of the data on the disk drive can take one or more hours, depending on system activity.

Note – All Sun StorEdge T3+ array disk drives either in a workgroup or an enterprise configuration must have the same storage capacity and drive speed.

Note – In some cases, if disk drives are removed and replaced on a powered-down or incompletely booted array, the internal system monitoring facilities may not correctly detect the drive replacement. Therefore, replace disk drives on powered-up, fully-booted arrays only.

Disk Drive LEDs

The top of each disk drive has LEDs that indicate drive activity and status. For information about the disk drive LEDs and how to interpret them to diagnose a drive failure, see Section , “Disk Drive LEDs” on page 13-156.

Note – Even if the LED indicates a drive failure, you should always verify the FRU status using either the CLI or Sun StorEdge Component Manager before replacing the drive. Refer to the *Sun StorEdge T3+ Array Administrator’s Manual* or the *Sun StorEdge Component Manager User’s Guide* for instructions.

▼ To Remove and Replace a Disk Drive

1. Observe static electricity precautions.
2. Remove the front panel by pressing in on the side latches (shown in FIGURE 14-1) and pulling the front panel forward.




FIGURE 14-1 Removing the Front Panel

3. Locate the disk drive that needs to be replaced.

Disk drives are numbered from 1 to 9, starting on the left side of the array (as shown in FIGURE 14-2).




FIGURE 14-2 Disk Drive Numbering

4. Use a coin or small screwdriver to press in and release the drive latch handle (FIGURE 14-3).




FIGURE 14-3 Releasing the Latch Handle

5. Use the latch handle to slowly pull the disk drive out an inch.

Wait 30 seconds before pulling the drive out completely, giving the disk drive time to spin down.

6. Remove the disk drive from the array (FIGURE 14-4).



Caution – Any disk drive that is removed must be replaced within 30 minutes or the array and all attached arrays will automatically shut down and power off.




FIGURE 14-4 Removing a Disk Drive

7. Push in the latch handle on the removed disk drive to protect it from damage.
8. Release the latch handle on the disk drive to be installed.
9. Insert the new disk drive gently on the middle of the rails and push it in until it is seated with the centerplane connector.
10. Use a coin or small screwdriver to press in and lock the latch handle.
11. Replace the front panel.



Caution – You must replace the front panel for the array to meet regulatory emissions requirements.

12. Using Sun StorEdge Component Manager or the CLI, verify that there are no disk drive errors.

Refer to the *Sun StorEdge Component Manager User's Guide* or the *Sun StorEdge T3+ Array Administrator's Manual* for more information.

Power and Cooling Units

The power and cooling unit has two active power sources: *standby* and *primary power*. Standby power, which is used to power the microcontroller on the interconnect card, is activated when AC power is present. Primary power, which is used to power all remaining circuits and disk drives, is activated when AC or battery power is present and the power switch is on.

Each power and cooling unit has a power switch in the rear upper center of the unit. Turning off the power on a power and cooling unit affects only that power and cooling unit. Therefore, to power off all primary power to the unit, both power switches on both power and cooling units must be turned off. After the switches are turned off, system primary power does not actually turn off until the controller has performed an orderly shutdown, including writing any data cache to disk. This process can take up to two minutes.

Separate power cords are used for the connector on each power and cooling unit to provide redundant cabling. The power cords need to be connected to alternate AC power sources for full redundancy (see FIGURE 14-5).




FIGURE 14-5 Power Cords Connected to the Power and Cooling Units



Caution – Do not handle the power and cooling unit when the power cord is connected. Line voltages are present within the power and cooling unit when the power cord is connected—even if the power switch is off.



Caution – At the rear of the power and cooling unit is a recessed PC card connector. Do not touch this connector or allow any metal object to touch it. The power and cooling unit contains the UPS battery backup. There are no serviceable parts inside this unit.

Power and Cooling Unit LEDs

Each power and cooling unit has an AC LED and a power supply (PS) LED. For information about the power and cooling unit LEDs and how to interpret them to diagnose a failure, see “Power and Cooling Unit LEDs” on page 171.

Note – Even if the LED indicates a power and cooling unit failure, you should always verify the FRU status using either the CLI or Sun StorEdge Component Manager before replacing the unit. Refer to the *Sun StorEdge T3+ Array Administrator’s Manual* or *Sun StorEdge Component Manager User’s Guide* for more information.

▼ To Remove and Replace a Power and Cooling Unit



Caution – To ensure correct airflow for system cooling, both power and cooling units must be in the installed position for normal operation. A failed power and cooling unit should be removed only when a replacement power and cooling unit is available to be inserted.



Caution – Replace only one power and cooling unit at a time to prevent system interruption.

To replace a power and cooling unit:

1. **Power off the power and cooling unit by pressing the power switch** (FIGURE 14-5). Make sure that the AC LED is amber and the PS LED is off.
2. **Disconnect the power cord from the AC outlet.**
3. **Disconnect the power cord from the power and cooling unit connector by squeezing both sides of the connector and pulling it straight out** (FIGURE 14-5).
4. **Unlock the power and cooling unit by using a coin or small screwdriver to push in and release the two latch handles** (FIGURE 14-6).
5. **Pull the power and cooling unit out of the array.**

Put one index finger through each of the latch handles. With your thumbs on the top of the chassis for support, pry the power and cooling unit out of its connectors with an upward rotation. After it is out approximately .5 inches (2 cm), it will be free to slide out of the frame on its rails.



Caution – Any power and cooling unit that is removed must be replaced within 30 minutes or the array and all attached arrays will automatically shut down and power off.




FIGURE 14-6 Removing the Power and Cooling Unit

- 6. Insert the new power and cooling unit.**
- 7. Lock the new power and cooling unit by pushing in both latch handles.**
- 8. Insert the power cord into the power and cooling unit connector.**
- 9. Connect the power cord to the AC outlet.**

Verify that the AC LED on the power and cooling unit is amber, indicating that AC power is present.
- 10. Push the power and cooling unit power switch on.**

Verify that both LEDs on the power and cooling unit are green, indicating that the unit is receiving power.

UPS Battery

The UPS battery is located within the power and cooling unit. The battery provides backup in case of a complete AC power failure and will sustain power to the array long enough to flush cache data to the drives.

During a power failure, if the battery is flushing cache data to the drives, battery power becomes depleted. After AC power is available, the battery recharges. While the battery is recharging, write-behind cache mode is disabled and write-through cache mode is enabled until the battery is fully recharged. The battery recharge could take up to 12 hours, depending on the length of the power outage and the amount of cache data that was flushed to the drives.

Battery Maintenance

The array has been factory configured to initiate a battery refresh cycle that occurs automatically once every 28 days. This battery refresh cycle involves a discharge of the battery followed by a recharge. The refresh cycle is sequential; the second power and cooling unit does not start a refresh until the first power and cooling unit successfully completes the refresh cycle, ensuring that only one battery in an array is refreshed at a time. This refresh cycle verifies the health of the battery. During the refresh, if a problem is detected with the battery, future refresh operations are suspended until the problem is fixed. If this happens, write-behind caching is turned off automatically as a safety precaution. A refresh cycle does not begin unless both power and cooling units are in optimal condition.

A battery refresh operation in progress is indicated in the `syslog` file. During the battery refresh operation, the battery health is determined by the discharge time of the battery. This time is written into the `syslog` file and is displayed as *battery hold time*. If the discharge time of the battery is greater than six minutes (360 seconds), the battery is considered healthy. If the battery is not healthy, the system will disable the battery, disable write cache, and indicate that the power and cooling unit needs to be replaced. The following is an example of the battery hold time as displayed in the `syslog` file:

```
May 21 00:19:22 driver24 BATD[1]: N: ulpcu1: hold time was 1159 seconds.
```

You can use the `refresh -s` command to view an active refresh operation. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* for more information on this command.

The battery service life is two years. When the battery approaches the end of its life span, warning messages are sent to the `syslog` file. The first message is sent 45 days before the end of the life span, followed by a warning message every 5 days thereafter. Replace the power and cooling unit that contains the battery as soon you receive these warning messages.

You can view information on the battery, including the life span, using the CLI or Sun StorEdge Component Manager. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* or the *Sun StorEdge Component Manager User's Guide* for instructions.

Removing and Replacing the UPS Battery

The UPS battery itself is *not* a serviceable FRU. To replace the battery, you must replace the entire power and cooling unit. The warning message will indicate which power and cooling unit needs to be replaced. See Section t, “To Remove and Replace a Power and Cooling Unit” on page 14-171 for instructions.

Interconnect Cards

The interconnect cards contain the interface circuitry and two connectors on each card for interconnecting multiple Sun StorEdge T3+ arrays. The interconnect card also contains the environmental monitor for the array. If an interconnect card or cable fails, the array continues to function using the remaining interconnect card and cable.

The interconnect cards in an array are assigned a *master* and *slave* status, similar to the controller card. If you unplug or remove the master interconnect card from the unit and then put it back, the new status is reported as “slave.” The corresponding interconnect card that was the slave then becomes the master. You can view the status of the interconnect cards using the `fru stat` command on the array. There is no performance impact to the array if the interconnect cards change status, and it is not necessary to alter the state back to the original master and slave configuration.

For more information on the `fru stat` command and other array commands, refer to the *Sun StorEdge T3+ Array Administrator's Manual*.

Interconnect Card LEDs

Each interconnect card has an interconnect status LED for each interconnect cable. For the location of and information about the interconnect card status LEDs and how to interpret them to diagnose a failure, see “Interconnect Card LEDs” on page 159.

Note – Even if the LED indicates an interconnect card failure, you should always verify the FRU status using either the CLI or Sun StorEdge Component Manager before replacing the unit. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* or *Sun StorEdge Component Manager User's Guide* for more information.

▼ To Removing and Replacing an Interconnect Card



Caution – Use the interconnect cables to cable only Sun StorEdge T3+ arrays together using the interconnect card connectors. Do *not* use these cables for any other FC-AL connection.



Caution – The interconnect card FRUs are extremely sensitive to static electricity. Use proper grounding wrist straps and antistatic procedures when handling these FRUs. See Section , “Preparing for Service” on page 14-165 for a list of static electricity precautions.



Caution – Replace one interconnect card at a time. Follow the procedure as described to ensure that there is no interruption in system operation or loss of data.

To prevent interruption of system operation during interconnect card replacement, make sure that:

- In a workgroup configuration, you remove only the failed interconnect card. Leave the second interconnect card intact in the array.



Caution – Removing both interconnect cards at the same time in a single unit will cause the unit to shut down immediately. Remove only one card at a time.

- In an enterprise configuration, you remove the interconnect cable only from the failed interconnect card. Leave the interconnect cable attached to the working interconnect card.

To replace an interconnect card:

1. **Ensure that the interconnect card to be replaced is showing failure status.**
2. **Observe static electricity precautions, as described in “Preparing for Service” on page 165.**

If you have a single controller unit configuration, proceed to Step 4.

3. **Remove the interconnect cable from the failed interconnect card only** (FIGURE 14-7).
Mark the connector with either 1 or 2.
4. **Unlock the failed interconnect card by pushing in on the latch handle.**




FIGURE 14-7 Removing the Interconnect Card

5. **Pull the interconnect card out using the latch handle.**



Caution – The interconnect card that is removed must be replaced within 30 minutes or the array and all attached arrays will automatically shut down and power off.

6. **Insert the new interconnect card, making sure that the card sits on the frame.**
7. **Lock the new interconnect card in place by pushing in on the latch handle.**
8. **Reconnect the interconnect cable to the interconnect card.**
9. **Verify that the LEDs on the interconnect card show that the interconnect card has initialized properly.**

You can verify the status of the interconnect card using either the CLI or Sun StorEdge Component Manager. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* or the *Sun StorEdge Component Manager User's Guide* for instructions.

Controller Card

The controller card is redundant only when the array is configured in an enterprise configuration; that is, when the array is connected to another array that contains a controller card.

In an enterprise configuration, if there is a controller card failure in a master unit, the controller would failover to the alternate master, which would then assume master controller status. After the failed controller card is replaced in the master unit, that unit assumes the role of alternate master. The new master unit (previously the alternate master) retains all the necessary information, such as the unit IP address, to continue functioning without system interruption. The enterprise configuration can be returned back to its original, prefailover configuration of master and alternate master by performing a `reset` command.

In a workgroup configuration, system activity is interrupted until the controller card is replaced and is back online. When a controller card is replaced, the array automatically downloads the previous settings to the new controller card. The array is fully functional after the new controller card is installed and the controller is online.

Controller Card LEDs

Note – The controller card LEDs for a Sun StorEdge T3 array differ from that of a Sun StorEdge T3+ array.

Note – Even if the LED indicates a controller card failure, you should always verify the FRU status using either the CLI or Sun StorEdge Component Manager before replacing the unit. Refer to the *Sun StorEdge T3+ Array Administrator's Manual* or *Sun StorEdge Component Manager User's Guide* for more information.

▼ To Remove and Replace a Controller Card

A controller card can be replaced without system interruption only if the array is part of an enterprise configuration.

To replace the controller card:

1. **Observe static electricity precautions.**
2. **Ensure that the controller card is showing failure status.**
3. **Remove the Ethernet cable from the Ethernet connector.**

4. Remove the fiber-optic cable from the FC-AL connector.
5. Unlock the controller card by pushing in on the latch handles (FIGURE 14-8).
6. Pull the controller card out using the latch handles.



Caution – The controller card that is removed must be replaced within 30 minutes or the array and all attached arrays will automatically shut down and power off.




FIGURE 14-8 Removing the Controller Card

7. Insert the new controller card.
8. Lock the new controller card by pushing in the latch handles.
9. Insert the fiber-optic cable into the FC-AL connector.
10. Insert the Ethernet cable into the Ethernet connector.
11. Check the controller status LED to determine when the controller is operational.

While the controller boots, the controller status LED is solid amber. When the controller is operational, the LED is green.

Note – After replacing the controller card, you might need to upgrade the controller firmware and EPROM firmware to make sure you are running the same firmware versions that you had on the previous controller card. To upgrade the controller firmware, see “Controller Firmware Version 2.1 Upgrade” on page 97.

Chassis

The array chassis rarely needs to be replaced. However, the chassis part number is available to replace the centerplane and chassis if necessary. (These must be replaced together, because they are factory aligned.)

Problems with the centerplane would most likely occur due to an electrical short or a bent or broken pin connector. These problems would first appear as a failure of another FRU component, such as an interconnect failure or drive failure. If replacing the FRU that appears to be failed does not correct the problem, then examine the centerplane connector that the FRU connects to for bent or broken pins. If nothing is obvious, install another spare FRU to verify that it is not a failed FRU component causing the problem. If you have eliminated all possibility of a FRU component failure and the problem still remains, it is likely to be a centerplane failure.

If there is a centerplane failure, contact your Sun Service representative. Only trained, qualified service providers should replace an array chassis and centerplane.

Specifications

Topics in this chapter include:

- “Power Specifications” on page 182
- “Environmental Specifications” on page 183
- “Mechanical Specifications” on page 184
- “Cable Specifications” on page 185

Power Specifications

Refer to TABLE A-1.

TABLE A-1 Power Specifications

| Specification | Rating |
|-----------------------|-----------------------------|
| Voltage and frequency | 100 to 240 VAC, 47 to 63 Hz |
| Input current | 5A max |
| Input power | 450W max |
| Heat load | 1540 BTUs/hr max |

Environmental Specifications

Refer to TABLE A-2.

TABLE A-2 Environmental Specifications

| Specifications | Operating | Non-operating | Notes |
|---|---|--|--|
| Temperature | 41°F (5°C) to 95°F (35°C) with maximum gradient 20°C per hour | 68°F (-20°C) to 140°F (60°C) with maximum gradient 20°C per hour | |
| Relative humidity | 20 to 80% nonconducting with maximum gradient 10% per hour | 5 to 93% noncondensing with maximum gradient 10% per hour | |
| Effective altitude (from drive specification) | -1,000 to +10,000 feet (-305 to +3,048 meters) | -1,000 to +40,000 feet (-305 to +12,192 meters) | |
| Shock (from drive specification) | 4.0G for maximum duration of 11 ms (half sinewave) | 10.0G for maximum duration of 11 ms (half sinewave) | Shock can be applied from any axis X,Y,Z |
| Vibration (from drive specification) | 5 to 500 Hz @ 0.25g (swept sinewave) | 5 to 500 Hz @ 1.0g (swept sinewave) | Vibration can be applied from any axis X,Y,Z |

Mechanical Specifications

To provide adequate clearance for FRU replacement, allow accessibility in front and rear of chassis. Allow a minimum of four inches in the front and rear for air circulation (the airflow direction is from front to rear). Refer to TABLE A-3.

TABLE A-3 Physical Specifications

| Height | Width | Depth | Weight |
|-------------|-------------|-------------|-----------------|
| 5.25 inches | 17.5 inches | 18.5 inches | 67 lbs |
| 13.33 cm | 44.45 cm | 46.99 cm | 30.39 kilograms |

Note – These specifications are for an individual array. For specifications on arrays in rack configurations, refer to the *Sun StorEdge T3+ Array Site Preparation Guide*.

Cable Specifications

Refer to TABLE A-4.

TABLE A-4 Cable Specifications

| Cable Type | Connector | Maximum Length |
|---|-------------------------------|-----------------------|
| Host interface (FC-AL): | | 82.021 ft (25 m) |
| • Sun StorEdge T3 array requires copper adapter (MIA) | DB-9 copper | |
| • Sun StorEdge T3+ array requires LC-SFF to SC fiber- optic cable | LC-SFF | |
| Shielded Ethernet: | RJ45 | 328.084 ft (100 m) |
| • 10/100BASE-T, category 5 (Sun StorEdge T3+ array) | | |
| • 10BASE-T, category 3 (Sun StorEdge T3 array) | RJ45 | |
| Serial: | RJ11 (Sun StorEdge T3 array) | 82.021 ft (25 m) |
| (Available only to qualified field service representatives.) | RJ45 (Sun StorEdge T3+ array) | |
| Power (110V) | Standard | n/a |
| Power (220V) | Standard | n/a |
| Array interconnect cable | DB-9 (not FC-AL compliant) | 2 ft (61 cm) |

Illustrated Parts

This appendix contains part numbers and illustrations of field-replaceable units (FRUs) for the Sun StorEdge T3+ array. See TABLE B-1.

TABLE B-1 Illustrated Part Numbers

| Part | Part Number | Page Number |
|---|--|-------------|
| Disk drive in canister; various capacities available | Contact your SunService representative | Page B-188 |
| Sun StorEdge T3+ array controller card | F501-5710 | Page B-189 |
| Interconnect card | F375-0085 | Page B-189 |
| Power and cooling unit | F300-1454 | Page B-189 |
| Interconnect cable: | | Page B-190 |
| • 14 in. (36 cm) | F530-2842 | |
| • 24 in. (61 cm) | F530-2843 | |
| Power cord | F180-1918 | Page B-190 |
| Sun StorEdge T3+ array 5m-fibre-optic cable | F537-1033 | Page B-190 |
| Sun StorEdge T3+ array shielded Ethernet cable (category 5) | F530-2991 | Page B-190 |
| Door assembly (front panel) | F540-4306 | Page B-191 |
| Empty array chassis with centerplane | F370-3990 | Not shown |




FIGURE B-1 Disk Drive in Canister




FIGURE B-2 Sun StorEdge T3+ Array Controller Card




FIGURE B-3 Interconnect Card




FIGURE B-4 Power and Cooling Unit




FIGURE B-5 Cables, Adapters, and Cords




FIGURE B-6 Door Assembly (Front Panel)

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 use 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 characteristic 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 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 very dependable.

**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 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

- AC power cord
 - connection, 68
- administration path, 148
- alternate master controller unit, 57
- Alternate Pathing (AP)
 - configuration recommendations, 18
- array
 - architecture of, 11
 - components of, 4
 - controller card and ports, 6
 - default configuration, 81
 - features of, 9
 - inspecting, 55
 - JumpStart, 57
 - overview of, 1
 - powering on, 70
 - verifying the firmware level, 75
- array disk tray, 11
 - cache modes, 129
 - commands, 102
 - components, 4
 - default settings, 94
 - description, 2
 - features, 9
 - FTP connection, 100
 - illustrated parts, 18, 191
 - installation, 93
 - LEDs, 156
 - operation, 133
 - overview of, 15
 - telnet connection, 101

auto cache mode, 116

B

- battery, 173
 - removing and replacing *See* power and cooling unit

C

- cabled partner group, 70
- cabling overview, 148
- cabling specifications, 185
- cabling the array, 65
- cache
 - allocation, 118
 - configuring, 116
 - for improving performance, 116
 - mirrored, enabling, 118
 - setting cache modes, 116
- cache modes, 129
- cache segment, 117
- capacity per unit, 10
- centerplane
 - service, 179
- changing the default configuration, 88
- channel connection failures, 161
- chassis
 - service, 179
- cluster support

See SunCluster 2.2 support

- components, 4
 - controller card description, 5
 - disk drive description, 4
 - power and cooling unit description, 9
- configurations
 - direct host
 - single host with eight controllers, 33, 34
 - single host with four controllers, 31
 - single host with one controller, 28
 - single host with two controllers, 29
 - enterprise, 16
 - hot spare, 121
 - hub host
 - dual hosts with eight controllers, 4, 43, 46, 47
 - dual hosts with four controllers, 40 to 41, 44, 45
 - single host with eight controllers, 39
 - single host with eight controllers>, 38
 - single host with four controllers, 35, 37
 - restrictions and recommendations, 17
 - scenarios, 27
 - Sun Cluster, 19
 - switch host
 - dual hosts with eight controllers, 50, 51
 - dual hosts with two controllers, 48, 49
 - workgroup, 14
- connecting the cables, 65
 - AC power cord, 68
 - Ethernet cable, 67
 - interconnect cables, 69
- connecting the host system, 89
- connections
 - Ethernet, 6
 - FC-AL, 6
 - serial, 6
- controller card, 5
 - array controller, 6
 - description, 5
 - illustrated, 189
 - LEDs, 160
 - removing and replacing, 177
 - service, 177
- controller failover, 91
- creating soft partitions, 90

D

- data, 86
- data block size
 - definition, 117
- data path, 148
- default configuration, 81
 - changing, 88
- disabling and reconstructing disk drives, 131
- disk drive, 4
 - description, 4
 - illustrated, 188
 - LEDs, 156
 - numbering, 81
 - removing and replacing, 166
 - status messages, 83
 - system area, 12
- disk tray settings, 94
- Dynamic Multi-Pathing (DMP)
 - configuration recommendations, 18

E

- editing the host files, 57
- enterprise configuration
 - configuration rules, 151
- environmental specifications, 183
- establishing a network connection, 72
- establishing logical volumes, 90
- Ethernet
 - administration path, 148
 - connection, 6
- Ethernet cable connection, 67

F

- FC-AL
 - data path, 148
- features
 - administration, 10
 - capacity per unit, 10
 - performance, 10
 - RAID capabilities, 10
 - reliability and redundancy, 10
 - scalability, 10
- fiber-optic cable

- illustrated, 190
- field-replaceable units (FRUs), 191
- front panel, 56
 - illustrated, 191
 - removing and replacing, 167
- ftp, 100

H

HBA

- SOC+, 142
- Sun StorEdge PCI Dual Fibre Channel network adapter, 144
- Sun StorEdge PCI FC-100, 141
- Sun StorEdge PCI Single Fibre Channel network adapter, 143
- Sun StorEdge SBus FC-100, 142

host files

- editing, 57

- host-generated messages, 155

hot spare

- default value, 125
- determining whether to use, 121

I

I/O boards

- Sun Enterprise SBus+ and Graphics+, 140

- illustrated parts, 187 to 191

- inspecting the array, 55

- installation, 93

- alternate pathing, 91
- checklist, 61
- configuring enterprise, 85
- configuring workgroup, 83
- connecting cables, 65
- connecting host system, 89
- connecting the host system, 89
- creating soft partitions, 90
- default configuration, 88
- editing the host files, 57
- establishing a network connection, 72
- establishing logical volumes, 90
- host files, 57
- hw configuration, 70
- mounting volumes, 81

- network connection, 72
- post instructions, 115
- power on/off, 94
- preparing, 53, 54
- required patches, 59
- setting the IP address, 57
- steps, 64
- verifying fw, 75
- verifying the firmware level, 76

interconnect cable

- connection, 69
- illustrated, 190

interconnect card

- description of, 7
- illustrated, 189
- in partner groups, 148
- LEDs, 159
- removing and replacing, 175
- service, 174

- IP address, assigning, 57

L

LEDs

- controller card, 160
- disk drive, 156
- interconnect card, 159
- power and cooling unit, 157

logical unit (LUN)

- See LUNs*

logical volumes

- post-installation configuration, 119

LUNs

- and applications, 120
- creating and labeling, 122
- default value, 125
- definition, 119
- determining how many are needed, 120
- guidelines for configuring, 119
- reconstruction rate, setting, 122
- viewing, 119

M

- MAC address, 17

- location of, 56

- master controller unit, 57
 - in a partner group, 137
 - parameters controlled by, 18
- mechanical specifications, 184
- media access control (MAC) address
 - See MAC address*
- MIA
 - illustrated, 190
- mounting volumes, 81

N

- network adapter
 - See HBA*
- network connection, 72

P

- parameters, tailored to I/O workload, 18
- partner group
 - alternate master unit, 57
 - configuration rules, 151
 - creating, 138
 - description, 14
 - direct host
 - single host with eight controllers, 33, 34
 - single host with four controllers, 31
 - single host with two controllers, 29
 - fully cabled, 70
 - how they work, 137
 - hub host
 - dual hosts with eight controllers, 46, 47
 - dual hosts with four controllers, 44, 45
 - single host with eight controllers, 38, 39
 - single host with four controllers, 35, 37
 - master unit, 57
 - multipathing software, 137
 - sharing parameter settings, 18
 - using multipathing software, 30
- performance, 10
- post-installation
 - array cabling, 147
 - array monitoring, 132
 - array settings, 127
 - concept of partner group, 136
 - configuring cache, 116

- configuring logical volumes, 119
- global parameters configuration, 115
- host connection, 139
- partner group configuration, 135
- RAID levels, 123
- power and cooling unit, 9
 - description, 9
 - illustrated, 189
 - LEDs, 157
 - removing and replacing, 171
- power cord
 - illustrated, 190
- power specifications, 182
- powering on, 70
- preparing for service, 165
- product
 - description of, 2
 - features of, 9
 - new features, 9

R

- RAID
 - and applications, 121
 - capabilities, 10
 - configuring for redundancy, 123
 - controller, 5
 - default level, 125
 - defined, 124
 - levels of, 125, 129
 - post-installation configuration, 123
- reliability and redundancy, 10
- removing and replacing
 - controller card, 177
 - disk drives, 166
 - front panel, 55, 167
 - interconnect card, 175
 - power and cooling unit, 171
 - UPS battery *See power and cooling unit*
- root password, 72

S

- scalability, 10
- serial number location, 56

- service
 - centerplane, 179
 - chassis, 179
 - controller card, 177
 - interconnect card, 174
 - power and cooling units, 170
 - preparing, 165
 - removing and replacing components, 165
- setting the IP address, 57
- SNMP notification, 132
- SOC+ HBA, 142
- software supported
 - clustering, 22
 - list of, 26
 - multi-pathing software, 22
 - multiple operational environment, 22
 - operational environment, 22
- specifications
 - cable, 185
 - environmental, 183
 - mechanical, 184
 - power, 182
- starting a telnet session, 72
- stripe unit size
 - See data block size
- Sun Cluster 2.2 support, 19
- Sun Enterprise SBus+ and Graphics+ I/O boards
 - See I/O boards
- Sun StorEdge PCI Dual Fibre Channel network adapter, 144
- Sun StorEdge PCI FC-100 HBA, 141
- Sun StorEdge PCI Single Fibre Channel network adapter, 143
- Sun StorEdge SBus FC-100 HBA, 142
- Sun StorEdge Traffic Manager
 - in configuration recommendations, 18
 - ipartner group configuration, 30
- supported software
 - bootability, 25
 - diagnostics and monitoring software, 23
 - host side drives, 23
 - import interface, 24
 - interfaces, 23
 - multi-host environments, 25
 - multi-platform, 25
 - SAN switch support, 23
- syslog command, 133

system area, 12

T

- telnet
 - command, 101
 - connection, 101
 - session, 72
- troubleshooting, 161
 - analyzing the problem, 155
 - channel connection failures, 161
 - FRU
 - failures, 161
 - telnet sessions, 155

U

- Unit interconnect card, 7
- unit volume configuration options, 128
- unpacking the array, 55
- upgrade
 - connecting to array, 100
 - limitations, 99
 - overview of, 98
 - requirements, 98
 - system firmware, 104
- UPS battery, 173

V

- verifying
 - firmware level, 75
 - hardware configuration, 70
- vol command
 - adding a volume, 86
 - initializing volumes, 86, 87
 - mounting volume, 87
- volume
 - initializing, 86, 87
 - mounting, 87

W

- workgroup configuration, 28

configuration rules, 150
write-behind cache mode, 116
write-through cache mode, 116