



Sun StorEdge™ 3000 Family Installation, Operation, and Service Manual

Sun StorEdge 3510 FC Array

Sun Microsystems, Inc.
www.sun.com

Part No. 816-7300-15
March 2004, Revision A

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

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

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

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

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

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

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

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

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

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

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

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

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

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



Please
Recycle



Adobe PostScript

Contents

Preface xxi

- 1. Product and Architecture Overview 1-1**
 - 1.1 Introduction 1-1
 - 1.2 Field-Replaceable Units (FRUs) 1-3
 - 1.2.1 RAID I/O Controller Modules 1-3
 - 1.2.2 I/O Expansion Modules 1-4
 - 1.2.3 Disk Drives 1-4
 - 1.2.4 Battery Module 1-5
 - 1.2.5 Power and Fan Modules 1-5
 - 1.3 Interoperability 1-5
 - 1.4 Fibre Channel Technology Overview 1-6
 - 1.4.1 FC Protocols 1-6
 - 1.4.2 FC Topologies 1-7
 - 1.4.3 Fibre Hubs and Switches 1-7
 - 1.4.4 Data Availability 1-7
 - 1.4.5 Scalability 1-8
 - 1.5 Fibre Channel Architecture 1-8
 - 1.5.1 Host and Drive FC Architecture 1-8
 - 1.5.2 Disk Drive FC Architecture 1-9

- 1.5.3 Redundant Configuration Considerations 1–10
 - 1.5.3.1 Host Bus Adapters 1–10
 - 1.5.3.2 Active-to-Active Redundant Controller 1–10
 - 1.5.3.3 Host Redundant Paths 1–10
- 1.6 Additional Software Tools 1–11
- 2. Site Planning 2–1**
 - 2.1 Customer Obligations 2–2
 - 2.2 Safety Precautions 2–2
 - 2.3 Environmental Requirements 2–3
 - 2.3.1 Electromagnetic Compatibility (EMC) 2–3
 - 2.4 Electrical and Power Specifications 2–4
 - 2.5 Physical Specifications 2–5
 - 2.6 Layout Map 2–5
 - 2.6.1 Rack Placement 2–5
 - 2.6.2 Tabletop Placement 2–6
 - 2.7 Console and Other Requirements 2–7
 - 2.8 Preinstallation Worksheet 2–7
- 3. Unpacking Your FC Array 3–1**
 - 3.1 Opening Your Package 3–2
 - 3.2 Checking the Package Contents 3–3
 - 3.2.1 Standard Sun StorEdge 3510 FC Array Package 3–3
 - 3.2.2 Field-Replaceable Units 3–4
 - 3.3 Customer-Provided Cables 3–5
 - 3.4 Mounting Your Array in a Rack or Cabinet 3–5
 - 3.5 Converting a JBOD to a RAID Array 3–5

- 4. Connecting Your Fibre Channel Array 4-1**
 - 4.1 Converting Your Front Bezel Locks So the Keys Cannot Be Removed 4-2
 - 4.2 Fibre Channel Array Connections 4-5
 - 4.3 Connecting the Chassis to an AC Power Outlet 4-6
 - 4.4 Connecting the Chassis to DC Power Outlets 4-7
 - 4.5 Powering Up and Checking LEDs 4-9
 - 4.6 Reviewing and Changing SFP Ports (Optional) 4-10
 - 4.6.1 Port Connectivity in a Dual-Controller Array 4-10
 - 4.6.2 Default SFP Placement 4-11
 - 4.6.3 Changing Your SFP Configuration 4-13
 - 4.7 Configuring a COM Port to Connect to a RAID Array 4-14
 - 4.8 Setting an IP Address 4-14
 - 4.9 Setting Up Out-of-Band Management Over Ethernet 4-16
 - 4.10 Cabling to Expansion Units 4-17
 - 4.10.1 Scaling Beyond 36 Disks 4-19
 - Note – Limitations 4-19
 - 4.11 Setting Loop IDs on Expansion Units 4-20
 - 4.12 Connecting Ports to Hosts 4-23
 - 4.13 Power-On Sequence 4-23
- 5. First-Time Configuration 5-1**
 - 5.1 Controller Defaults and Limitations 5-2
 - 5.1.1 Planning for Reliability, Availability, and Serviceability 5-2
 - 5.1.2 Dual-Controller Considerations 5-3
 - 5.1.3 Single-Controller Considerations 5-4
 - 5.2 Battery Operation 5-4
 - 5.2.1 Battery Status 5-4
 - 5.2.2 Write-Back and Write-Through Cache Options 5-6
 - 5.3 Software Management Tools 5-6

- 5.3.1 Out-of-Band Connection 5–6
- 5.4 Configuration Overview 5–7
 - 5.4.1 Point-to-Point Configuration Guidelines 5–9
 - 5.4.2 A SAN Point-to-Point Configuration Example 5–11
 - 5.4.3 A DAS Loop Configuration Example 5–15
 - 5.4.4 Larger Configurations 5–18
- 5.5 Initial Configuration Steps 5–18
 - 5.5.1 Viewing the Initial Firmware Windows 5–19
 - 5.5.2 Checking Available Physical Drives 5–20
 - 5.5.3 Configuring FC Channels as Host or Drive (Optional) 5–22
 - 5.5.4 Choosing Loop or Point-to-Point Fibre Connection 5–24
 - 5.5.5 Editing and Creating Additional Host IDs (Optional) 5–25
 - 5.5.6 Selecting Sequential or Random Optimization 5–27
 - 5.5.6.1 Maximum Number of Disks and Maximum Usable Capacity for Random and Sequential Optimization 5–28
 - 5.5.7 Reviewing Default Logical Drives and RAID Levels 5–29
 - 5.5.8 Completing Basic Configuration 5–31
 - 5.5.9 Creating Logical Drives (Optional) 5–31
 - 5.5.10 Preparing for Logical Drives Larger Than 253 Gbyte 5–38
 - 5.5.11 Changing a Logical Drive Controller Assignment (Optional) 5–39
 - 5.5.12 Creating or Changing a Logical Drive Name (Optional) 5–41
 - 5.5.13 Partitioning a Logical Drive (Optional) 5–42
- 5.6 Mapping Logical Drive Partitions to Host LUNs 5–45
 - 5.6.1 Planning for 1024 LUNs (Optional, Loop Mode Only) 5–47
 - 5.6.2 First Steps in Mapping a Partition to a LUN 5–48
 - 5.6.3 Using the Map Host LUN Option 5–50
 - 5.6.4 Setting Up Host Filter Entries 5–52
 - 5.6.4.1 Creating Host Filter Entries 5–53

- 5.6.5 Creating Device Files for the Solaris Operating Environment 5–58
- 5.6.6 Saving Configuration (NVRAM) to a Disk 5–59
- 5.7 Using Software to Monitor and Manage the Sun StorEdge 3510 FC Array 5–60
 - 5.7.1 Other Supported Software 5–60
 - 5.7.2 Enabling VERITAS DMP 5–61
 - 5.7.3 The VERITAS Volume Manager ASL 5–61
- 6. Checking LEDs 6–1**
 - 6.1 LEDs When Array Is First Powered On 6–1
 - 6.2 Front Panel LEDs 6–1
 - 6.2.1 Correcting SES or PLD Firmware Version Conflicts 6–4
 - 6.3 Back Panel LEDs 6–5
 - 6.3.1 I/O Controller Module LEDs 6–5
 - 6.3.2 Power Supply and Fan Module LEDs 6–8
- 7. Maintaining Your Array 7–1**
 - 7.1 Introducing Key Screens and Commands 7–2
 - 7.1.1 Controller Firmware Initial Screen 7–2
 - 7.1.2 Main Menu 7–4
 - 7.1.3 Quick Installation (Reserved) 7–4
 - 7.2 Silencing Audible Alarms 7–5
 - 7.3 Deleting Logical Drives 7–6
 - 7.4 Checking Status Windows 7–8
 - 7.4.1 Logical Drive Status Table 7–9
 - 7.4.2 Physical Drive Status Table 7–11
 - 7.4.3 Channel Status Table 7–13
 - 7.4.4 Controller Voltage Temperature Status 7–15
 - 7.4.5 Viewing SES Status 7–18
 - 7.4.6 SES Temperature Sensor Locations (FC Only) 7–20

- 7.4.7 Identifying Fans 7-21
- 7.4.8 Viewing Event Logs on the Screen 7-22
- 7.5 Restoring Your Configuration (NVRAM) From a File 7-24
- 7.6 Upgrading Firmware 7-26
 - 7.6.1 Patch Downloads 7-27
 - 7.6.2 Controller Firmware Upgrade Features 7-27
 - 7.6.3 Installing Firmware Upgrades 7-28
 - 7.6.4 Installing Controller Firmware Upgrades From the Firmware Application 7-29
- 7.7 Replacing the Front Bezel and Ear Caps 7-30
 - 7.7.1 Removing the Front Bezel and Ear Caps 7-30
 - 7.7.2 Placing the Bezel and Ear Caps Back Onto the Chassis 7-31
- 8. Troubleshooting Your Array 8-1**
 - 8.1 RAID LUNs Not Visible to the Host 8-1
 - 8.2 Controller Failover 8-2
 - 8.3 Rebuilding Logical Drives 8-3
 - 8.3.1 Automatic Logical Drive Rebuild 8-3
 - 8.3.2 Manual Rebuild 8-5
 - 8.3.3 Concurrent Rebuild in RAID 1+0 8-6
 - 8.4 Identifying a Failed Drive for Replacement 8-7
 - 8.4.1 Flash Selected Drive 8-9
 - 8.4.2 Flash All SCSI Drives 8-10
 - 8.4.3 Flash All But Selected Drive 8-10
 - 8.5 Recovering From Fatal Drive Failure 8-11
 - 8.6 SES Temperature Sensor Locations 8-13
 - 8.7 Identifying Fans From the SES Device Status Menu 8-14
 - 8.8 Modifying Drive-Side SCSI Parameters 8-15
 - 8.9 Troubleshooting Flowcharts 8-15

8.9.1	Power Supply and Fan Module	8-16
8.9.2	Drive LEDs	8-19
8.9.3	Front Panel LEDs	8-21
8.9.4	I/O Controller Module	8-25
8.10	Using the Reset Button	8-27
8.11	Silencing Audible Alarms	8-27
A. Sun StorEdge 3510 FC Array Specifications A-1		
A.1	Physical Specifications	A-2
A.2	Sun StorEdge 3510 FC Array Highlights	A-3
A.2.1	Hardware Loop ID	A-3
A.2.2	Firmware Host-Side Connection Mode	A-4
A.2.3	Firmware LUN Filtering (RAID-Based Mapping)	A-4
A.2.4	Firmware Redundant Loops	A-4
A.2.5	Firmware Dynamic Load Balancing	A-4
A.3	Agency Approvals and Standards	A-5
B. Using Standalone JBOD Arrays B-1		
B.1	Supported Configurations (JBOD Arrays)	B-2
B.2	Supported Operating Systems (JBOD Arrays)	B-2
B.3	Supported Host Platforms and Connection Methods (JBOD Arrays)	B-2
B.4	Known Limitations Affecting Sun StorEdge 3510 FC JBOD Arrays	B-3
B.5	Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays	B-4
B.5.1	Sun StorEdge Configuration Service	B-4
B.5.1.1	Enabling JBOD Support	B-5
B.5.2	Sun StorEdge Diagnostic Reporter	B-6
B.5.3	Sun StorEdge CLI	B-7
B.6	Setting Loop IDs on Sun StorEdge 3510 FC JBOD Arrays	B-7
B.6.1	Changing Your ID Switch Setting	B-8

- B.7 Connecting Sun StorEdge 3510 FC JBOD Arrays B-9
 - B.7.1 Single-Port Connections to Sun StorEdge 3510 FC JBOD Arrays B-9
 - B.7.2 Dual-Port Connections to Sun StorEdge 3510 FC JBOD Arrays B-10
 - B.7.3 Changing Your SFP Configuration on JBOD Arrays B-12
 - B.7.4 Connecting the Sun StorEdge 3510 FC JBOD Array to a Host Computer B-13
- B.8 Enabling Sun StorEdge 3510 FC JBOD Array Support B-13
- B.9 Replacing a Disk Drive in a Sun StorEdge 3510 FC JBOD Array B-15
 - B.9.1 Removing a Disk Drive B-16
 - B.9.2 Installing a New Disk Drive B-16
 - B.9.3 Installing an Air Management Sled B-16
- B.10 Downloading Firmware to Disk Drives in a JBOD B-17
- B.11 Understanding Component Status LEDs on JBOD Arrays B-17
 - B.11.1 LED Activity When JBOD Arrays Are Powered On B-17
 - B.11.2 Front Panel LEDs on JBOD Arrays B-17
 - B.11.3 Back Panel LEDs on JBOD Arrays B-19
 - B.11.4 Back Panel I/O Module LEDs B-20
 - B.11.5 Back Panel Power Supply and Fan Module LEDs on JBOD Arrays B-21
- B.12 Troubleshooting Sun StorEdge 3510 FC JBOD Arrays B-22
 - B.12.1 Troubleshooting Configuration Issues B-22
 - B.12.2 Troubleshooting Hardware Issues B-23
- B.13 Converting JBODs to RAID arrays B-27

C. Failed Component Alarm Codes C-1

D. Record of Settings D-1

- D.1 View and Edit Logical Drives D-2
 - D.1.1 Logical Drive Information D-2
 - D.1.2 Logical Drive Partition Information D-3
- D.2 View and Edit Host LUNs D-4
 - D.2.1 LUN Mappings D-4
- D.3 View and Edit SCSI Drives D-5
- D.4 View and Edit SCSI Channels D-6
- D.5 View and Edit Peripheral Devices D-7
 - D.5.1 View System Information D-7
- D.6 Save NVRAM to Disk and Restore From Disk D-7

E. Cable Pinouts E-1

- E.1 RJ-45 Connector E-1
- E.2 DB-9 COM port E-2

F. Configuring a Sun Server Running the Solaris Operating Environment F-1

- F.1 Setting Up the Serial Port Connection F-1
- F.2 Accessing the Firmware Application From a Solaris Host F-2
- F.3 Using the `tip` Command for Local Access to the Array F-3
- F.4 Determining the WWN in the Solaris Operating Environment F-3

G. Configuring a Windows 2000 Server or Windows 2000 Advanced Server G-1

- G.1 Setting Up the Serial Port Connection G-2
- G.2 Accessing the Firmware Application From a Windows 2000 Server or Windows 2000 Advanced Server G-5
- G.3 Enabling a Windows 2000 Server or Windows 2000 Advanced Server to Recognize New Devices and LUNs G-5
- G.4 Determining the World Wide Name for Windows 2000 Servers and Windows 2000 Advanced Servers G-10

H. Configuring a Linux Server H-1

- H.1 Setting Up the Serial Port Connection H-1
- H.2 Accessing the Firmware Application From a Linux Server H-3
- H.3 Checking the Adapter BIOS H-4
- H.4 Multiple LUN Linux Configuration H-5
- H.5 Making an ext3 File System for Linux H-6
- H.6 Creating a File System H-7
- H.7 Creating a Mount Point and Mounting the File System Manually H-7
- H.8 Mounting the File System Automatically H-8
- H.9 Determining the World Wide Name for Linux Hosts H-9

I. Configuring an IBM Server Running the AIX Operating Environment I-1

- I.1 Setting Up a Serial Port Connection I-2
- I.2 Accessing the Firmware Application From an IBM Server Running AIX I-3
- I.3 Identifying the Device on Which You Will Create a Logical Volume I-4
- I.4 Using SMIT to Enable an AIX Host to Recognize New LUNs I-5
- I.5 Creating a Volume Group I-6
- I.6 Creating a Logical Volume I-7
- I.7 Creating a File System I-7
- I.8 Mounting the New File System I-8
- I.9 Verifying That the New File System Is Mounted I-9
- I.10 Determining the World Wide Name for IBM Servers Running AIX I-9

J. Configuring an HP Server Running the HP-UX Operating Environment J-1

- J.1 Setting Up a Serial Port Connection J-2
- J.2 Accessing the Firmware Application From an HP Server Running HP-UX J-3
- J.3 Attaching the Disk Array J-5
- J.4 Logical Volume Manager J-6

- J.5 Definitions of Common Terms J-6
- J.6 Creating a Physical Volume J-7
- J.7 Creating a Volume Group J-7
- J.8 Creating a Logical Volume J-9
- J.9 Creating an HP-UX File System J-10
- J.10 Mounting the File System Manually J-10
- J.11 Mounting the File System Automatically J-10
- J.12 Determining the World Wide Name for HP-UX Hosts J-11

K. Configuring a Windows NT Server K-1

- K.1 Setting Up the Serial Port Connection K-2
- K.2 Accessing the Firmware Application From a Windows NT Server K-5
- K.3 Enabling a Windows NT Server to Recognize New Devices and LUNs K-5
- K.4 Determining the World Wide Name for Windows NT Servers K-9

Index Index-1

Figures

- [FIGURE 1-1](#) Sun StorEdge 3510 FC Array Front View 1–1
- [FIGURE 4-1](#) Front Bezel and Front Bezel Locks of an Array 4–2
- [FIGURE 4-2](#) Sequence of Steps to Change Front Bezel Locks So Keys Cannot Be Removed 4–3
- [FIGURE 4-3](#) Hardware Connections on the Back of a Dual-Controller FC Array 4–5
- [FIGURE 4-4](#) Installed Cord Lock 4–7
- [FIGURE 4-5](#) Front Panel With LEDs Displayed 4–9
- [FIGURE 4-6](#) Dedicated Drive Channels 2 (on the Upper Controller) and 3 (on the Lower Controller) in a Dual-Controller Array 4–11
- [FIGURE 4-7](#) Host Channels on a Dual-Controller Array 4–11
- [FIGURE 4-8](#) Default Dual-Controller SFP Placement 4–12
- [FIGURE 4-9](#) Default Single-Controller SFP Placement 4–12
- [FIGURE 4-10](#) Default JBOD Expansion Unit SFP Placement 4–13
- [FIGURE 4-11](#) Typical SFP Connector Used to Connect Cables to Chassis SFP Ports 4–13
- [FIGURE 4-12](#) RAID Array Attached to Two Hosts and Two Expansion Units 4–18
- [FIGURE 4-13](#) Front Bezel and Front Bezel Locks of an Array 4–20
- [FIGURE 4-14](#) ID Switch Located on the Left Front Side of the Array or Expansion Unit 4–21
- [FIGURE 5-1](#) A Point-to-Point Configuration With a Dual-Controller Array and Two Switches 5–13
- [FIGURE 5-2](#) A DAS Configuration With Four Servers, a Dual-Controller Array, and Two Expansion Units 5–16
- [FIGURE 5-3](#) Firmware Main Menu 5–20
- [FIGURE 5-4](#) Example of an Allocation of Local and Spare Drives in Logical Configurations 5–32

FIGURE 5-5	Partitions in Logical Configurations	5-42
FIGURE 5-6	Mapping Partitions to Host ID/LUNs	5-47
FIGURE 5-7	Example of LUN Filtering	5-52
FIGURE 6-1	Front Panel LEDs	6-2
FIGURE 6-2	Chassis Ear LEDs and Reset Button on Front Panel	6-2
FIGURE 6-3	I/O Controller Module and Battery Module LEDs	6-5
FIGURE 6-4	I/O Expansion Module for an Expansion Unit	6-6
FIGURE 6-5	AC Power Supply and Fan Module	6-8
FIGURE 6-6	DC Power Supply and Fan Module	6-8
FIGURE 7-1	Cooling Fan Locations	7-22
FIGURE 8-1	Automatic Rebuild	8-4
FIGURE 8-2	Manual Rebuild	8-6
FIGURE 8-3	Flashing the Drive LED of a Selected Drive	8-9
FIGURE 8-4	Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive	8-10
FIGURE 8-5	Flashing All Drive LEDs Except a Selected Drive LED	8-11
FIGURE 8-6	Cooling Fan Locations	8-15
FIGURE 8-1	Power Supply or Fan Module Flowchart, 1 of 2	8-17
FIGURE 8-2	Power Supply or Fan Module Flowchart, 2 of 2	8-18
FIGURE 8-3	FC Drive LEDs Flowchart, 1 of 2	8-20
FIGURE 8-4	FC Drive LEDs Flowchart, 2 of 2	8-21
FIGURE 8-5	Front Panel LEDs (FC) Flowchart, 1 of 4	8-22
FIGURE 8-6	Front Panel LEDs Flowchart, 2 of 4	8-23
FIGURE 8-7	Front Panel LEDs Flowchart, 3 of 4	8-24
FIGURE 8-8	Front Panel LEDs Flowchart, 4 of 4	8-25
FIGURE 8-9	I/O Controller Module Flowchart	8-26
FIGURE B-1	ID Switch	B-7
FIGURE B-2	Sun StorEdge 3510 FC JBOD Array Attached to a Single HBA Port	B-10
FIGURE B-3	Sun StorEdge 3510 FC JBOD Array Attached to Two HBA Ports	B-11
FIGURE B-4	Front View of a Drive Module Pulled Out of the Chassis	B-15
FIGURE B-5	Front Panel LEDs	B-18

FIGURE B-6	Chassis Ear LEDs and Reset Button on Front Panel	B-18
FIGURE B-7	I/O Module for a Sun StorEdge 3510 FC JBOD Array	B-20
FIGURE B-8	AC Power Supply and Fan Module	B-21
FIGURE B-9	DC Power Supply and Fan Module	B-21
FIGURE B-10	JBOD or Expansion Unit Flowchart 1 of 2	B-25
FIGURE B-11	JBOD or Expansion Unit Flowchart 2 of 2	B-26
FIGURE E-1	Ethernet RJ-45 Socket 10/100 BASE-T	E-1
FIGURE E-2	RS-232 DB9 (EIA/TIA 574) View of the Male End	E-2
FIGURE F-1	RAID Array COM Port Connected Locally to the Serial Port of a Workstation	F-2
FIGURE F-2	Worldwide Name Information Displayed by the <code>luxadm</code> command.	F-4
FIGURE I-1	RAID Array COM Port Connected Locally to the Serial Port of a Host System	I-3
FIGURE I-2	Network Address Corresponding to WWN	I-10
FIGURE J-1	RAID Array COM Port Connected Locally to the Serial Port of a Host System	J-3

Tables

TABLE 1-1	Sun StorEdge 3510 FC Array Configuration Options	1–2
TABLE 2-1	Environmental Specifications	2–3
TABLE 2-2	Power Specifications	2–4
TABLE 2-3	Physical Specifications	2–5
TABLE 2-4	Preinstallation Worksheet	2–8
TABLE 2-5	Host and Fabric Switch Connectivity Summarized	2–9
TABLE 3-1	List of Available FRUs	3–4
TABLE 4-1	DC Cable Wiring for Cable 35-00000148	4–8
TABLE 4-2	DC Cable Wiring for Cable 35-00000156	4–8
TABLE 4-3	ID Switch Settings for Expansion Units	4–22
TABLE 4-4	Sample Array and Expansion Units With Different Loop IDs and Drive IDs	4–22
TABLE 5-1	Battery Status Indicators	5–5
TABLE 5-2	Example Point-to-Point Configuration With Two Logical Drives in a Dual-Controller Array	5–14
TABLE 5-3	Example Primary and Secondary ID Numbers in a Loop Configuration With Two IDs per Channel	5–17
TABLE 5-4	Default Host Channel IDs	5–25
TABLE 5-5	Maximum Number of Disks per Logical Drive for a 2U Array	5–28
TABLE 5-6	Maximum Usable Capacity (Gbyte) per Logical Drive for a 2U Array	5–29
TABLE 5-7	RAID Level Definitions	5–30
TABLE 5-8	Cylinder and Head Mapping for the Solaris Operating Environment	5–38

TABLE 5-9	Cylinder and Head Mapping for the Solaris Operating Environment	5–39
TABLE 5-10	Configuration for 1024 LUNs	5–48
TABLE 6-1	Front Panel LED Status When Array Is First Powered On	6–1
TABLE 6-2	Front Panel LEDs	6–3
TABLE 6-3	I/O Controller Module and Battery Module LEDs	6–6
TABLE 6-4	Power Supply LEDs	6–8
TABLE 7-1	Components of the Controller Firmware Screen	7–2
TABLE 7-2	Progress Indicator Prefix Meanings	7–3
TABLE 7-3	Silencing the Alarm	7–6
TABLE 7-4	Parameters Displayed in the Logical Drive Status Window	7–9
TABLE 7-5	Parameters Displayed in the Drive Status Window	7–11
TABLE 7-6	Parameters Displayed in the Channel Window	7–14
TABLE 7-7	Sun StorEdge 3510 Temperature Sensor Locations	7–21
TABLE 7-8	Relationship Between Cooling Elements, Fans, and Power Supply Modules	7–22
TABLE 8-1	Sun StorEdge 3510 Temperature Sensor Locations	8–13
TABLE 8-2	Relationship Between Cooling Elements, Fans, and Power Supply Modules	8–14
TABLE A-1	Sun StorEdge 3510 FC Array Physical Specifications	A–2
TABLE B-1	Supported Sun Servers and Connection Methods (JBOD Arrays)	B–2
TABLE B-2	ID Switch Settings for Sun StorEdge 3510 FC JBOD Arrays	B–8
TABLE B-3	Example of 12 Drives Shown on a 12-Disk Sun StorEdge 3510 FC JBOD Array Connected to a Host Over a Single FC Loop	B–10
TABLE B-4	Example of 24 Drives Shown on a 12-Disk Sun StorEdge 3510 FC JBOD Array Connected to a Host Over Two FC Loops	B–12
TABLE B-5	Front Panel LED Status When Array Is First Powered On	B–17
TABLE B-6	JBOD Front Panel LEDs	B–19
TABLE B-7	JBOD I/O Controller Module and Battery Module LEDs	B–20
TABLE B-8	Power Supply LEDs	B–22
TABLE C-1	Failed Component Alarm Codes	C–1
TABLE E-1	Ethernet RJ-45 Pin Description	E–1
TABLE E-2	Pin Names	E–2

Preface

This manual provides instructions for installing, initially configuring, and operating your Sun StorEdge™ 3510 FC Array.

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



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

How This Book Is Organized

This book covers the following topics:

[Chapter 1](#) provides an overview of the array's features.

[Chapter 2](#) covers site planning and basic safety requirements.

[Chapter 3](#) provides general procedures for unpacking and inspecting the array.

[Chapter 4](#) provides procedures for connecting your array to power and to the network.

[Chapter 5](#) provides procedures for initially configuring your array.

[Chapter 6](#) describes the array's front and back panel LEDs.

[Chapter 7](#) describes maintenance procedures.

[Chapter 8](#) describes troubleshooting procedures.

[Appendix A](#) provides Sun StorEdge 3510 array specifications.

[Appendix B](#) provides information about standalone JBOD arrays.

[Appendix C](#) provides information about failed component alarm codes.

[Appendix D](#) provides tables for recording configuration data.

[Appendix E](#) provides pinout identification for each connector.

[Appendix F](#) provides instructions on configuring a Solaris™ server.

[Appendix G](#) provides instructions on configuring a Windows 2000 or Windows 2000 Advanced server.

[Appendix H](#) provides instructions on configuring a Linux server.

[Appendix I](#) provides instructions on configuring a IBM AIX server.

[Appendix J](#) provides instructions on configuring an HP-UX server.

[Appendix K](#) provides instructions on configuring a Windows NT server.

Using UNIX Commands

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

- Software documentation that you received with your system
- Solaris™ operating environment documentation, which is at

<http://docs.sun.com>

Shell Prompts

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

Typographic Conventions

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

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

Related Documentation

Title	Part Number
<i>Sun StorEdge 3510 FC Array Release Notes</i>	816-7301
<i>Sun StorEdge 3000 Family Best Practices Manual</i>	816-7325
<i>Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide</i>	817-3711
<i>Sun StorEdge 3000 Family Configuration Service 1.5 User's Guide</i>	817-3337
<i>Sun StorEdge 3000 Family Diagnostic Reporter 1.5 User's Guide</i>	817-3338
<i>Sun StorEdge 3000 Family 1.5 Software Installation Manual</i>	817-3764
<i>Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays</i>	817-3629
<i>Sun StorEdge 3000 Family CLI 1.5 User's Guide</i>	817-4951
<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326
<i>Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual</i>	816-7930

Accessing Sun Documentation

All Sun StorEdge 3510 FC Array documentation is available online in both PDF and HTML format at the following locations:

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

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

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

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

You can order printed copies of manuals for the Sun StorEdge 3510 FC Array at:

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

Contacting Sun Technical Support

For late-breaking news and troubleshooting tips, review the *Sun StorEdge 3510 FC Array Release Notes* located at:

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

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

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

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

800-USA-4SUN

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

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

508 Accessibility Features

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

Sun Welcomes Your Comments

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

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

Please include the title and part number of your document with your feedback:

Sun StorEdge 3000 Family Installation, Operation, and Service Manual, part number 816-7300-15.

Product and Architecture Overview

This chapter provides a brief overview of your Sun StorEdge 3510 FC (Fibre Channel) Array. Topics covered in this chapter are:

- “Introduction” on page 1-1
- “Field-Replaceable Units (FRUs)” on page 1-3
- “Interoperability” on page 1-5
- “Fibre Channel Technology Overview” on page 1-6
- “Fibre Channel Architecture” on page 1-8
- “Additional Software Tools” on page 1-11

1.1 Introduction

The Sun StorEdge 3510 FC Array is a rack-mountable, Network Equipment Building System (NEBS) Level 3-compliant, Fibre Channel mass storage subsystem. NEBS Level 3 is the highest level of NEBS criteria used to assure maximum operability of networking equipment in mission-critical environments such as telecommunications central offices. The array is designed for high availability, high performance, and high capacity.



FIGURE 1-1 Sun StorEdge 3510 FC Array Front View

The Sun StorEdge 3510 FC Array models include:

- A redundant array of independent disks (RAID) with a single controller
- A RAID array with two controllers
- An expansion unit that connects to and is managed by a RAID array
- A Just a Bunch of Disks (JBOD) array that has no RAID controller and connects to, and is managed by, a host server

See [“Using Standalone JBOD Arrays” on page B-1](#) for detailed information about using JBOD arrays.

Note – A label on the bottom lip of an array chassis, underneath the front bezel, indicates whether it is a JBOD array or a RAID array. For instance, “3510 AC JBOD” refers to an alternating-current version of a JBOD array, “3510 DC JBOD” refers to a direct-current version of a JBOD array, and “3510 AC RAID” refers to an alternating-current version of a RAID array. Similarly, using a UNIX command such as `probe-scsi-all` provides similar information, using an “A” designator for RAID arrays and a “D” designator for disks in a JBOD array. For example, “StorEdge 3510F D1000” identifies a JBOD array with SES firmware version 1000.

[TABLE 1-1](#) shows the configuration options for the Sun StorEdge 3510 FC Array.

TABLE 1-1 Sun StorEdge 3510 FC Array Configuration Options

Internal RAID controllers	Up to 2, with a minimum of 1
2 GHz / 1 GHz Fibre Channel disks	Up to 12 per array or per expansion unit, with a minimum of 4 plus 1 spare
Fibre Channel expansion units*	Up to 8
Fibre Channel JBOD arrays†	1
Connection options	<ul style="list-style-type: none"> • Serial port • Ethernet • Fibre Channel Small Form-Factor Pluggable (SFP)
Supported RAID levels	0, 1, 3, 5, 1+0, 3+0, and 5+0
Redundant field-replaceable units (FRUs)	<ul style="list-style-type: none"> • Power supply and fan modules • I/O controller modules • I/O expansion modules • Battery board module • Disk drive modules
Configuration management and enclosure event reporting options‡	<ul style="list-style-type: none"> • In-band Fibre Channel • Out-of-band 10/100 BASE-T Ethernet • RS-232 connectivity • Enclosure monitoring by SCSI Enclosure Services (SES)

- * A disk array with no controller. Each expansion unit has two Fibre Channel loops that can provide redundant data paths back to the RAID array.
- † A disk array with no controller that is connected directly to a host computer, with no RAID array in the loop.
- ‡ The host-based Sun StorEdge Configuration Service software provides a graphical user interface (GUI) and additional event-reporting capabilities.

For a list of supported racks and cabinets, refer to the release notes for the Sun StorEdge 3510 FC array. You can find these release notes at:

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

Reliability, availability, and serviceability (RAS) are supported by:

- Redundant components
- Notification of failed components
- Components that are replaceable while the unit is online

For information about specifications and agency approvals, see “Sun StorEdge 3510 FC Array Specifications” on page A-1.

1.2 Field-Replaceable Units (FRUs)

This section describes the FRUs contained in the Sun StorEdge 3510 FC Array.

1.2.1 RAID I/O Controller Modules

A dual-controller configuration offers increased reliability and availability because it eliminates a single point of failure, the controller. In a dual-controller configuration, if the primary controller fails, the array automatically fails over to the second controller without an interruption of data flow.

The Sun StorEdge 3510 FC Array I/O controller modules are hot-swappable. Each of these RAID controller modules provides six Fibre Channel ports that support 2-gigabit (Gb) or 1-Gb data rates. Single- and dual-controller models are available, with the dual-controller version supporting active/passive and active/active configurations. Each 3510 RAID controller is configured with 1 gigabyte (Gbyte) of cache.

In the unlikely event of a failure, a redundant RAID controller immediately begins servicing all I/O requests. The failure does not affect application programs.

Each RAID controller module can support up to 1 gigabyte of Synchronous Dynamic Random Access Memory (SDRAM) with Error Control Check (ECC) memory. In addition, each controller supports 64 megabytes (Mbyte) of on-board memory. Two Application Specific Integrated Circuit (ASIC) controller chips handle the interconnection between the controller bus, DRAM memory, and Peripheral Component Interconnect (PCI) internal buses. They also handle the interface between the on-board 2 Mbyte flash, 32 Kbyte nonvolatile random access memory (NVRAM) RS-232 port chip, and 10/100 BASE-T Ethernet chip.

The RAID controller module is a multifunction board that provides six Small Form-Factor Pluggable (SFP) ports, the SES logic, and the RAID controller. The SES logic monitors various temperature thresholds, fan speed from each fan, voltage status from each power supply, and the FRU ID.

Each RAID controller module incorporates SES direct-attached Fibre Channel capability to monitor and maintain enclosure environmental information. The SES controller chip monitors all internal +12 and +5 voltages, various temperature sensors located throughout the chassis, each fan, and controls the front and back panel LEDs and audible alarm. Both the RAID chassis and the expansion chassis support dual SES failover capabilities for fully redundant event monitoring.

1.2.2 I/O Expansion Modules

The hot-serviceable I/O expansion modules also provide four SFP ports but do not have battery modules or controllers, and are used only in expansion units.

1.2.3 Disk Drives

Each disk drive is mounted in its own sled assembly, known as an *air management sled*. Each sled assembly has EMI shielding, an insertion and locking mechanism, and a compression spring for maximum shock and vibration protection.

The drives can be ordered in 36-GB, 73-GB, and 146-GB sizes. Thirty-six-GB drives have a rotation speed of 15,000 RPM, while 73-GB drives and 146-GB drives have a rotation speed of 10,000 RPM. Each disk drive is slot-independent, meaning that once a RAID set has been initialized, the system can be shut down and the drives can be removed and replaced in any order. In addition, disk drives are field-upgradeable to larger drives without interruption of service to user applications. The drive firmware is also field-upgradeable, but the firmware upgrade procedure requires interruption of service.

In the event of a single disk drive failure, with the exception of RAID 0, the system continues to service all I/O requests. Either the mirrored data or parity data is utilized in the rebuilding of the failed drives data to a spare disk drive, assuming one is assigned. If a spare is not assigned, you have to manually rebuild the array.

In the unlikely event that multiple disk drive failures occur within the same RAID set, data that has not been replicated or backed up might be lost. This is an inherent limitation of all RAID subsystems and could affect application programs.

An air management sled FRU is available for use when you remove a disk drive and do not replace it. You can insert the air management sled into the empty slot to maintain optimum airflow through the chassis.

1.2.4 Battery Module

The battery module is designed to provide power to the system cache for 72 hours in the case of a power failure. When power is reapplied, the cache is purged to disk. The battery module is a hot-swappable FRU that is mounted on the I/O board with guide rails and a transition board. It also contains the EIA-232 and DB-9 serial interface (COM) ports.

1.2.5 Power and Fan Modules

Each array contains redundant (two) power and fan modules. Each module contains a 420-watt power supply with autoranging capabilities from 90 Volts Alternating Current (VAC) to 264 VAC and two radial 52 Cubic Feet per Minute (CFM) fans. A single power and fan module can sustain an array.

1.3 Interoperability

The array is designed for heterogeneous operation and supports the following operating environments:

- Solaris operating environment versions 8 and 9
- Sun™ Linux 5.0 on the Sun LX50 server
- Red Hat Linux distribution
- Windows NT 4.0 and Windows 2000 server
- IBM AIX Terminal Menu
- HP-UX

Note – For information about supported versions of these operating environments, refer to the release notes for your array.

The array does not require any host-based software for configuration, management, and monitoring, which are handled through the built-in firmware application. The console window can be accessed via the DB-9 communications (COM) port using the `tip` command, or via the Ethernet port using the `telnet` command.

1.4 Fibre Channel Technology Overview

As a device protocol capable of high data transfer rates, Fibre Channel simplifies data bus sharing and supports not only greater speed than SCSI, but also more devices on the same bus. Fibre Channel can be used over both copper wire and optical cable. It can be used for concurrent communications among multiple workstations, servers, storage systems, and other peripherals using SCSI and IP protocols. When a Fibre Channel hub or Fabric switch is employed, it provides flexible topologies for interconnections.

1.4.1 FC Protocols

Two common protocols are used to connect Fibre Channel (FC) nodes together:

- Point-to-point
- Arbitrated loop

The point-to-point protocol is straightforward, doing little more than establishing a permanent communication link between two ports.

The arbitrated loop protocol creates a simple network featuring distributed (arbitrated) management between two or more ports using a circular (loop) data path. Arbitrated loops can support more nodes than point-to-point connections.

The Sun StorEdge 3510 FC Array supports point-to-point and arbitrated loop protocols. You select the protocol you prefer by setting the desired FC connection option in the configuration parameters of the firmware application.

1.4.2 FC Topologies

The presence or lack of switches establishes the topology of an FC environment. In a direct attached storage (DAS) topology, servers connect straight to arrays without switches. In a storage area network (SAN) topology, servers and arrays connect to an FC network created and managed by switches.

Refer to *Sun StorEdge 3000 Family Best Practices Manual* for your array to see information about optimal configurations for site requirements.

1.4.3 Fibre Hubs and Switches

A storage network built on a Fibre Channel architecture might employ several of the following components: Fibre Channel host adapters, hubs, Fabric switches, and fibre-to-SCSI bridges.

- **Fibre hubs.** An arbitrated loop hub is a wiring concentrator. “Arbitrated” means that all nodes communicating over this fibre loop are sharing a 100 megabits per second (Mbps) segment. Whenever more devices are added to a single segment, the bandwidth available to each node is further divided.

A loop configuration allows different devices in the loop to be configured in a token ring style. With a fibre hub, a fibre loop can be rearranged in a star-like configuration because the hub itself contains port bypass circuitry that forms an internal loop inside. Bypass circuits can automatically reconfigure the loop once a device is removed or added without disrupting the physical connection to other devices.

- **Fabric switches.** A Fabric switch functions as a routing engine, which actively directs data transfer from source to destination and arbitrates every connection. Bandwidth per node via a Fabric switch remains constant when more nodes are added, and a node on a switch port uses an up-to-100-Mbps data path to send or receive data.

1.4.4 Data Availability

Data availability is one of the major requirements for today’s mission-critical applications. Highest availability can be accomplished with the following functionality:

- **Hot-plug capabilities.** With proper hardware and software configuration in dual-controller mode, a failed controller can be replaced online while the existing controller is actively serving I/O.
- **Dual-loop configurations.** Dual loop provides path redundancy and greater throughput.

- **Controller communications over Fibre Channel.** Selectable either through dedicated loops or all drive loops. This allows a more flexible configuration of redundant controllers.

1.4.5 Scalability

The Fibre Channel architecture brings scalability and easier upgrades to storage. Storage expansion can be as easy as cascading another expansion unit to a configured RAID array without powering down the running system. Up to two expansion units can be daisy-chained and connected to a single RAID array.

Up to 125 devices can be configured in a single FC loop. By default, the array provides two drive loops and four host/drive loops, and operates in Fibre Channel-Arbitrated Loop (FC-AL) and Fabric topologies.

1.5 Fibre Channel Architecture

Each RAID array has six Fibre Channels that have the following defaults:

- Channels 0, 1, 4, and 5 are host channels connected to servers and can be reassigned as drive channels to connect to expansion units.
- Channels 2 and 3 are dedicated drive channels that connect the internal 12-disk drives in the RAID chassis, and can also be used to add expansion chassis to the configuration.
- FC-AL is the default mode. Point-to-point is also available.

The expansion unit has a total of four FC-AL ports.

Note – Throughout this manual, Fibre Channel-Arbitrated Loops are referred to simply as *loops*.

1.5.1 Host and Drive FC Architecture

RAID controller channels 0, 1, 4, and 5 are designated for host connections or for connections to expansion chassis drives, depending on controller configuration settings.

In a dual RAID controller configuration, both RAID controllers have the same channel designators, due to the architecture of the loops within the chassis. Each host and drive channel of the top RAID controller shares a loop with the matching channel on the bottom RAID controller. For example, Channel 0 of the top RAID controller shares the same loop as channel 0 of the bottom RAID controller. This provides four separate loops for direct connectivity to hosts, expansion chassis, or hub and switch devices.

Each host and drive loop includes multiple components. Each I/O board contains FC-AL loops and port bypass circuits that are associated with three components: a channel on the RAID controllers, the SFP port residing within the I/O board, and a connection to the opposite I/O board. This architecture provides a data path for both RAID controllers to either the top or bottom SFP port on a given channel. Plugging an SFP transceiver into a port enables external connections to that port. The optical connector used is the low-profile LC connector. Plugging and unplugging SFP connectors into ports is straightforward and easy.

The configuration of a single RAID controller configuration is slightly different. The connection to the lower I/O board does not exist; however, the same number of loops are available.

1.5.2 Disk Drive FC Architecture

Channels 2 and 3 are disk drive channels only and cannot be used as host channels. In the disk drive FC architecture, channel 2 from both controllers appears on the upper I/O board and channel 3 from both controllers appears on the lower I/O board.

The components residing on the disk drive loops are:

- Channel 2 or 3 (depending upon I/O board being reviewed) from one controller
- Channel 2 or 3 from the second controller
- Two SFPs
- Disk drive loop A for channel 2 or disk drive loop B for channel 3
- The SES logic

Note – RAID array channels 0, 1, 4, and 5 can be host or drive ports; by default, they are host ports. The CH2 and CH3 ports are dedicated drive ports and cannot be used as host ports.

1.5.3 Redundant Configuration Considerations

This section provides information about setting up redundant configurations for increased reliability.

1.5.3.1 Host Bus Adapters

Fibre Channel is widely applied to storage configurations with topologies that aim to avoid loss of data by component failure. As a rule, the connections between source and target should be configured in redundant pairs.

The recommended host-side connection consists of two or more host bus adapters (HBAs). Each HBA is used to configure a Fibre Channel loop between the host computer and the array. In active-to-active redundant controller mode, the primary loop serves the I/O traffic directed to the primary controller, and its pair loop serves the I/O traffic to the secondary controller. The host-side management software directs I/O traffic to the pair loop if one of the redundant loops fails.

1.5.3.2 Active-to-Active Redundant Controller

Since each fibre interface supports only a single loop ID, two HBAs are necessary for the active-to-active redundant controller operation. Using two HBAs in each server ensures continued operation even when a data path fails.

In active-to-active mode, the connection to each host adapter should be considered a data path connecting the host to either the primary or the secondary controller. One adapter should be configured to serve the primary controller and the other adapter to serve the secondary controller. Each target ID on the host channels should be assigned either a primary ID or a secondary ID. If one controller fails, the existing controller can inherit the ID from its counterpart and activate the one standby channel to serve host I/O.

1.5.3.3 Host Redundant Paths

The controller passively supports redundant fibre loops on the host side, provided that the host has implemented software support for this feature.

In the unlikely event of controller failure, the standby channels on an existing controller become an I/O route serving the host I/O originally directed to the active channel on its pair of controllers. Moreover, application failover software should be running on the host computer to control the transfer of I/O from one HBA to another in case either data path fails.

1.6 Additional Software Tools

The following additional software tools are available on the Sun StorEdge 3000 Family Professional Storage Manager CD, provided with your array:

- Sun StorEdge Configuration Service, a management and monitoring program
- Sun StorEdge Diagnostic Reporter software, a monitoring utility
- Sun StorEdge CLI, a command-line utility to monitor and manage the array.

Refer to the *Sun StorEdge 3000 3000 Family Software Installation Guide* for information about installing these tools.

The Sun StorEdge 3000 Family Documentation CD provides the related user guides with configuration procedures for these tools.

For other supported software tools, refer to the release notes, located at:

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

Site Planning

This chapter outlines the site-planning requirements and basic safety requirements for the installation and use of the Sun StorEdge 3510 FC Array. Complete a preinstallation worksheet and prepare the site for installation according to the worksheet details and the specified site-planning requirements.

Review the details of this chapter before installing a Sun StorEdge 3510 FC Array. Topics covered in this chapter are:

- “Customer Obligations” on page 2-2
- “Safety Precautions” on page 2-2
- “Environmental Requirements” on page 2-3
- “Electrical and Power Specifications” on page 2-4
- “Physical Specifications” on page 2-5
- “Layout Map” on page 2-5
- “Console and Other Requirements” on page 2-7
- “Preinstallation Worksheet” on page 2-7

Note – The release notes for your array list the supported operating environments, host platforms, software, and qualified cabinets.

2.1 Customer Obligations

The customer is obliged to inform Sun Microsystems, Inc. of any and all ordinances and regulations that would affect installation.



Caution – When selecting an installation site for the Sun StorEdge 3510 FC Array, choose a location that avoids excessive heat, direct sunlight, dust, or chemical exposure. Such exposure greatly reduces the product’s longevity and might void your warranty.

The customer is responsible for meeting all government codes and regulations concerning facilities. The customer is also responsible for compliance with the following requirements:

- Meeting all local, national, and international codes covered in this specification. The subjects covered include fire and safety, building, and electrical codes.
- Documenting and informing Sun Microsystems, Inc. of any deviations from “[Sun StorEdge 3510 FC Array Specifications](#)” on page A-1.

2.2 Safety Precautions

For your protection, observe the following safety precautions when setting up your equipment:

- Follow all safety precautions and requirements specified in the *Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual*.
- A fully loaded array weighs over 50 pounds. Use two people to lift the array to avoid injury.
- Follow all cautions and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source match the voltage and frequency inscribed on the equipment’s electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous voltages might be present. Conductive foreign objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

- To reduce the risk of electric shock, do not plug Sun products into any other type of power system. Sun products are designed to work with single-phase power systems having a grounded neutral conductor. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.
- Your Sun product is shipped with a grounding-type (three-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.
- Do not use household extension cords with your Sun product. Not all power cords have the same current ratings. Household extension cords do not have overload protection and are not meant for use with computer systems.
- Do not block or cover the openings of your Sun product. Never place a Sun product near a radiator or heat register. Failure to follow these guidelines can cause overheating and affect the reliability of your Sun product.

2.3 Environmental Requirements

TABLE 2-1 Environmental Specifications

	Operating	Nonoperating
Altitude	To 3000 meters (9000 feet)	To 12,000 meters (36,000 feet)
Humidity	10% to 90% RH at 27 degrees C max wet bulb (noncondensing)	93% RH, 38 degrees C max wet bulb temperature (noncondensing)
Temperature		
Standalone	5 degrees C to 40 degrees C	-40 degrees C to +65 degrees C
Rack	5 degrees C to 35 degrees C	-40 degrees C to +65 degrees C

2.3.1 Electromagnetic Compatibility (EMC)

The following is required for all installations:

- All AC mains and supply conductors to power distribution boxes for both the rack-mounted array and the desktop array must be enclosed in a metal conduit or raceway when specified by local, national, or other applicable government codes and regulations.
- The supply conductors and power distribution boxes (or equivalent metal enclosure) must be grounded at both ends.

- The supplied arrays require voltages within minimum fluctuation.
- The facilities voltage supplied by the customer must maintain a voltage with not more than (+/-) 5 percent fluctuation. The customer facilities must provide suitable surge protection.

2.4 Electrical and Power Specifications

The Sun StorEdge 3510 FC Array requires two independent power sources. Each array has two power supply and fan modules for redundancy.

Each AC array requires two 115 VAC/15A or two 240 VAC service outlets. All AC power supplies are autoranging and are automatically configured to a range of 90-264 VAC and 47-63 Hz. There is no need to make special adjustments.

Each DC array requires two -48 VDC service outlets and has an input voltage range of -36 VDC to -72 VDC.

Note – To ensure power redundancy, be sure to connect the two Sun StorEdge 3510 FC power modules to two separate circuits (for example, one commercial circuit and one UPS).

TABLE 2-2 Power Specifications

AC power	Voltage and frequency 90 to 264 VAC, 47 to 63 Hz
Input current	5A max
Power-supply output voltage	+5 VDC and +12 VDC
DC power	-48V DC (-36 VDC to -72 VDC)

2.5 Physical Specifications

Use the physical specifications in [TABLE 2-3](#) to plan the location of your array.

TABLE 2-3 Physical Specifications

Category	Description
Dimensions	2U (3.45 inches / 8.76 cm) height 20 inches / 50.8 cm chassis depth 17.5 inches / 44.6 cm width (19 inches / 48.26 cm with ears)
Installation clearances	For FRU component removal and replacement, 15-inches (37 cm) is required front and back.
Cooling clearances	6 inches (15 cm) is required front and back. No cooling clearance is required on the sides or the top and bottom of the array.

2.6 Layout Map

It is helpful to create a sketch or layout map to indicate the exact location for the Sun StorEdge 3510 FC Array as well as the location of the hosts, console, and Ethernet connections that will be connected to it.

As you lay out the components, consider the cable lengths that will be used.

2.6.1 Rack Placement

Follow these guidelines when preparing a rackmount placement for your system:

- Ensure that the floor surface is level.
- Leave enough space in front of the rack to access components for servicing.
- Leave enough space in back of the rack to access components for servicing.
- Keep power and interface cables clear of foot traffic. Route cables inside walls, under the floor, through the ceiling, or in protective channels or raceways.
- Route interface cables away from motors and other sources of magnetic or radio frequency interference.
- Stay within the cable length limitations.

- Provide two separate power sources for the array. These power sources must be independent of each other, and each must be controlled by a separate circuit breaker at the power distribution point.

2.6.2 Tabletop Placement

The Sun StorEdge 3510 FC Array can be positioned on a desk or a table. Follow these guidelines when preparing a tabletop placement for your system:

- Choose a desk or a table that can support 50 pounds for each fully configured array you plan to place on it.
- Do not place the arrays on the edge of the table. Set the array so that at least 50 percent of the array is inside the table or desk leg support area. Failure to do this might cause the table to tip over.
- Leave enough space in front and in back of the array to access components for servicing. To remove the components requires a clearance of 15 inches (37 cm) in front and in back of the array.
- Provide a minimum space of 6 inches (15 cm) in front and in back of the array for adequate airflow.
- Keep power and interface cables clear of foot traffic. Route cables inside walls, under the floor, through the ceiling, or in protective channels or raceways.
- Route interface cables away from motors and other sources of magnetic or radio frequency interference.
- Stay within the cable length limitations.
- Ensure that the operating environment for the array does not exceed the specifications.
- Use two people to lift the array to avoid injury. The array can weigh up to 50 pounds.
- Do not place the array in a vertical position. Place the array horizontally.
- If you are installing more than one array, you can stack up to five arrays on top of each other. Do not stack more than five arrays together.
- Provide two separate power sources for the array. These power sources must be independent of each other, and each must be controlled by a separate circuit breaker at the power distribution point.

2.7 Console and Other Requirements

A console with at least one serial port or Ethernet connection is necessary for installation and configuration of your Sun StorEdge 3510 FC Array. Once you have configured your array with an IP address, the Ethernet port can also be useful for configuring the array.

Refer to the following preinstallation worksheet for additional preparation details.

2.8 Preinstallation Worksheet

Before ordering a Sun StorEdge 3510 FC Array, complete the preinstallation worksheet and then prepare the site for installation according to the site-planning requirements.

Note – If you are connecting to several hosts or Fabric switches, make as many copies of [TABLE 2-5](#) as you need and label them appropriately.

You are responsible for ensuring that the site consistently conforms to all stipulated standards and that necessary peripherals are made available to the engineer during installation.

Review the details of your specific survey before installing your Sun StorEdge 3510 FC Array.

If necessary, attach or sketch a network diagram to the survey.

TABLE 2-4 Preinstallation Worksheet

Rackmounting	<p>Customers must ensure that the appropriate service outlets are available for installation. Requirements vary.</p> <p>Will the Sun StorEdge 3510 FC Array be rackmounted? Yes / No</p> <ul style="list-style-type: none">• Is the rack supplied by Sun Microsystems, Inc.? Yes / No• If yes, include Sun model number: _____• If not, make/model: _____ / _____ <p>Does the rack mount:</p> <ul style="list-style-type: none">• Front and back? If so, depth? _____• Center/Telco? _____ <p>What cable lengths are required? _____</p> <p>Are there any power strips or power sequencers in the rack? Yes / No</p> <p>Are they supplied by Sun Microsystems, Inc.? Yes / No If yes, part number: _____</p> <p>If not, quantity of plugs/outlets required: _____ / _____</p>
IP address	<p>Array IP address: _____.</p> <p>Array network mask: _____.</p>
Cabling	<p>FC cable lengths to connect to hosts: _____</p>

TABLE 2-5 Host and Fabric Switch Connectivity Summarized

Host or Fabric Switch Connectivity - Host or Fabric Switch #1

Host or Fabric switch name: _____

Host or Fabric switch make/model: _____

HBA connector types: _____

Cable distance from the array to the host(s): _____

Operating environment: _____

Patches installed: _____

IP addresses:

- Network _____
 - Host or switch _____
-

Host or Fabric Switch Connectivity - Host or Fabric Switch #2

Host or Fabric switch name: _____

Host or Fabric switch make/model: _____

HBA connector types: _____

Cable distance from the array to the host(s): _____

Operating environment: _____

Patches installed: _____

IP addresses:

- Network _____
 - Host or switch _____
-

Unpacking Your FC Array

This chapter describes the procedure for unpacking your Sun StorEdge 3510 FC Array package. The following topics are covered in this chapter:

- [“Opening Your Package” on page 3-2](#)
- [“Checking the Package Contents” on page 3-3](#)
- [“Customer-Provided Cables” on page 3-5](#)
- [“Mounting Your Array in a Rack or Cabinet” on page 3-5](#)
- [“Converting a JBOD to a RAID Array” on page 3-5](#)

3.1 Opening Your Package

Follow these guidelines for unpacking the equipment.



Caution – To avoid personal injury or damage to the equipment during installation, always use two people to remove the unit from its container. This unit weighs approximately 50 pounds.

1. **Select a suitable area for unpacking.**
2. **Store all packing material and boxes for possible equipment returns.**
3. **Check the contents sheet in your product package.**

The contents sheet summarizes the standard contents for your product. See [“Checking the Package Contents” on page 3-3](#) for more information.

4. **Compare the packing slip and the list of parts with the items you received.**

If the list of parts on your packing slip does not match the items you received, or any items appear damaged, immediately notify your carrier agent and the supplier who prepared your shipment.

5. **Carefully examine the cables provided in the package.**

If any cable appears to be damaged, contact the technical service department for an immediate replacement.

6. **Check the list of [“Customer-Provided Cables” on page 3-5](#).**

These are required to complete your installation.

Note – You must purchase or provide Fibre Channel cables for connecting the Sun StorEdge 3510 FC Array to host servers.

3.2 Checking the Package Contents

Inspect your Sun StorEdge 3510 FC Array packages for standard items as well as purchased options before you begin installation. If any parts are missing or damaged, contact your sales representative immediately.

3.2.1 Standard Sun StorEdge 3510 FC Array Package

Quantity	Item
1	One or more of the following arrays: <ul style="list-style-type: none">• Sun StorEdge 3510 FC Array with single controller• Sun StorEdge 3510 FC Array with dual controllers• Sun StorEdge 3510 JBOD Expansion Unit
1	<ul style="list-style-type: none">• <i>Sun StorEdge 3510 FC Array Contents Sheet</i>• To download and print the latest Sun StorEdge 3510 FC Array release notes, go to: http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Workgroup/3510
2	CDs in a CD set; one CD for software and one CD for user documentation
1	Serial null modem cable
1 or 2	25-foot (7.5m) CAT5 Ethernet cable per controller in each array
1	Cable adapter, DB9-to-DB25
2	DC power cables if you ordered a DC-powered array
2	AC cord locks in a plastic bag if you ordered an AC-powered array
2	Front bezel keys in a plastic bag, to secure the front bezel onto the chassis
Various	Purchased options. These options are ordered at the time of purchase and are integrated into or added to the unit prior to delivery.

3.2.2 Field-Replaceable Units

Check that you received all field-replaceable units (FRUs) ordered with your Sun StorEdge 3510 FC Array. For additional FRUs, consult your sales representative. For instructions on how to install or replace FRUS, review the following manuals located on the product web site:

- *Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays*
- *Sun StorEdge 3000 Family FRU Installation Guide*

TABLE 3-1 List of Available FRUs

FRU Model Number	Description
F370-5535-01	BOX, 2U, FC, Chassis + Backplane (RAID/JBOD)
F370-5545-01	BATTERY, FC, 2U
F370-5540-01	CABLE, FC, 1.5 FT, EXPANSION
F370-5537-01	I/O w/SES & RAID CONT FC, 1 GB MEMORY, BATTERY, 2U
F370-5538-01	I/O w/SES, JBOD FC, 2U
F370-5398-01	AC PWR SUPPLY/FAN MOD, 2U
XTA-3310-DC-Kit	DC PWR SUPPLY/FAN MOD, 2U
XTA-3510-36GB-15K	DRIVE MOD, 36 GB FC, 15KRPM
XTA-3510-73GB-10K	DRIVE MOD, 73 GB FC, 10KRPM
XTA-3510-146GB-10K	DRIVE MOD, 146 GB FC, 10KRPM
XTA-3510-Ctrl-1G	I/O w/SES & RAID CONT FC, 1 GB MEMORY, BATTERY, 2U
XTA-3000-AMBS	AIR MANAGEMENT SLED
XTA-3310-RK-19S	RACK KIT, 2U, 19" CABINET 22"-28"
XTA-3310-RK-19L	RACK KIT, 2U, 19" CABINET 28"-36"
XTA-3310-RK-19C	RACK KIT, 2U, 19" CABINET TELCO CENTER MNT
XTA-3310-RK-19F	RACK KIT, 2U, 19" CABINET TELCO FRONT MNT
XSFP-LW-2GB	SFP, 2G,LW 1310NM, FC, TRANS
XSFP-SW-2GB	SFP, 2G, SW 850NM,FC, TRANS

All FRUs are hot-swappable except the I/O controller and I/O expansion modules, which is hot-serviceable. *Hot-serviceable* means that the module can be replaced while the array and hosts are powered up but the connected hosts must be inactive.

For information about how to install FRUs, refer to the *Sun StorEdge 3000 Family FRU Installation Guide*.

3.3 Customer-Provided Cables

Customers must provide a minimum of one Fibre Channel cable per host to connect a host to a RAID array. Two Fibre Channel cables are required for a redundant path configuration. To obtain qualified cables, consult your Sun sales representative.

3.4 Mounting Your Array in a Rack or Cabinet

Refer to *Sun StorEdge 3000 Family Rack Installation Guide* for instructions on how to install and prepare a rack or cabinet for mounting your FC array.

3.5 Converting a JBOD to a RAID Array

If you want to convert a JBOD to a RAID array, refer to the *Sun StorEdge 3000 Family FRU Installation Guide*.

Connecting Your Fibre Channel Array

This chapter describes procedures for cabling the Sun StorEdge 3510 FC Array and for connecting the array to power and to network devices.

The following topics are covered in this chapter:

- [“Converting Your Front Bezel Locks So the Keys Cannot Be Removed” on page 4-2](#)
- [“Fibre Channel Array Connections” on page 4-5](#)
- [“Connecting the Chassis to an AC Power Outlet” on page 4-6](#)
- [“Connecting the Chassis to DC Power Outlets” on page 4-7](#)
- [“Powering Up and Checking LEDs” on page 4-9](#)
- [“Reviewing and Changing SFP Ports \(Optional\)” on page 4-10](#)
 - [“Port Connectivity in a Dual-Controller Array” on page 4-10](#)
 - [“Default SFP Placement” on page 4-11](#)
 - [“Changing Your SFP Configuration” on page 4-13](#)
- [“Configuring a COM Port to Connect to a RAID Array” on page 4-14](#)
- [“Setting an IP Address” on page 4-14](#)
- [“Setting Up Out-of-Band Management Over Ethernet” on page 4-16](#)
- [“Cabling to Expansion Units” on page 4-17](#)
- [“Setting Loop IDs on Expansion Units” on page 4-20](#)
- [“Connecting Ports to Hosts” on page 4-23](#)
- [“Power-On Sequence” on page 4-23](#)

Before you connect the Sun StorEdge 3510 FC Array to the network, position the Sun StorEdge 3510 FC Array in the rack, cabinet, or other location where it will be used.



Caution – When positioning the array, do not block the air vents at the front or back of the unit. Follow all safety precautions specified in the *Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual*.



Caution – When you power off an array, wait five seconds before you power it back on. If you power the array off and on too quickly, unexpected results can occur.

4.1 Converting Your Front Bezel Locks So the Keys Cannot Be Removed

The bezel on your array has two locks whose keys can be removed when the locks are in either the locked or open position. It is possible to reconfigure the locks so that the keys cannot be removed.




FIGURE 4-1 Front Bezel and Front Bezel Locks of an Array

To change the locks so the keys cannot be removed, follow these steps:

1. **Remove the bezel by gently pivoting the swing arms out of their ear sockets.**
See [“Removing the Front Bezel and Ear Caps”](#) on page 7-30 for step-by-step instructions on how to remove the bezel.
2. **Make sure the key is in the locked position, with the pawl extending horizontally past the edge of the bezel.**




FIGURE 4-2 Sequence of Steps to Change Front Bezel Locks So Keys Cannot Be Removed

3. Hold the key in place and use a 12 mm or 3/8-inch nut driver to remove the locking nut that holds the pawl in place, as shown in the first panel of [FIGURE 4-2](#).



Caution – Be sure to hold the key in place. Otherwise there is a risk of breaking the small tab on the lock that serves as a stop.

4. Lift the pawl off the threaded part of the lock body, as shown in the second panel of [FIGURE 4-2](#).
5. Set the pawl aside, face up, so that you can remember its orientation when you replace it.

6. Use the key to turn the lock 180 degrees, as shown in the third panel of [FIGURE 4-2](#)
7. Replace the pawl in the same orientation as before, as shown in the fourth panel of [FIGURE 4-2](#).
8. Hold the key in place and use the nut driver to refasten the locking nut that holds the pawl in place, as shown in the fifth panel of [FIGURE 4-2](#). Be careful not to cross-thread the nut.



Caution – Be sure to hold the key in place. Otherwise there is a risk of breaking the small tab on the lock that serves as a stop.

9. Replace the bezel.

Note – To convert your bezel locks so that the keys can be removed, repeat this procedure.

4.2 Fibre Channel Array Connections

FIGURE 4-3 identifies the hardware connections on the back of a dual-controller FC array.




FIGURE 4-3 Hardware Connections on the Back of a Dual-Controller FC Array

Management is in-band through fibre host connections and out-of-band through the serial port and Ethernet port on the back of each controller.

4.3 Connecting the Chassis to an AC Power Outlet

When you connect the AC power cords, you should install the provided two cord locks at the same time. The provided AC cord locks are used to securely fasten the AC cable connectors.



Caution – If the array is connected to AC power sources not within the designated 90–135, 180–265 VAC PFC range, damage might occur to the unit.

Note – To ensure power redundancy, be sure to connect the two power supply modules to two separate circuits (for example, one commercial circuit and one UPS).

To connect the AC power cords, perform the following procedure:

1. Obtain an appropriate AC power cable.
2. Use a screwdriver to remove the screw and cylindrical standoff from one of the two cord locks provided and set them aside for reassembly later.
3. Slide the cord lock over the AC power connector.
4. Hold the cylindrical standoff between the two screw-holes on the flanges of the cord lock.
5. Insert the screw into the first screw-hole, through the standoff, and then into the threaded screw-hole on the other flange.
6. Tighten the screw with a screwdriver until the flanges bottom out on the cylindrical standoff.
7. Push the power cord into the power supply receptacle until it is firmly seated.
8. Push the green ejector handle forward until it is seated against the power supply.
9. Turn the thumbscrew of the green ejector handle clockwise until it is finger-tight to secure the handle and the cord lock.




FIGURE 4-4 Installed Cord Lock

10. Repeat this procedure for the second cord lock and second power cable.

4.4 Connecting the Chassis to DC Power Outlets

Two DC power cords are packaged with each DC array. To connect the DC power cords, perform the following procedure:

1. Connect a DC power cable to the first power supply and to a power outlet.

Note – Use only the DC power cables provided with the array.

2. Check the DC cable part number and wire labels carefully before connecting the cable to the source.

TABLE 4-1 DC Cable Wiring for Cable 35-00000148

Pin Number	Voltage	Color
A3	Return	Red
A2	GND (Chassis Ground)	Green/Yellow
A1	-48vdc	Black

TABLE 4-2 DC Cable Wiring for Cable 35-00000156

Pin Number	Voltage	Color
A3	L+	Red
A2	GND (Chassis Ground)	Green/Yellow
A1	L-	White



Caution – If the Sun StorEdge 3510 FC Array is connected to DC power sources not within the designated –48V DC (–36 VDC to –72 VDC) range, damage might occur to the unit.

Note – To ensure power redundancy, be sure to connect the two power supply modules to two separate circuits (for example, one commercial circuit and one UPS).

Note – To extend the length of the DC power cable as needed, strip the last 1/4” of the cable, insert the stripped end into a provided Panduit tube, and crimp the tube.

3. Tighten the cable-locking screws to attach the cable securely to the power supply power outlet.
4. Connect the second power cable to the second power supply and to a second power outlet. Tighten the cable-locking screws.
If one power supply fails, the second power supply automatically takes the full load.

4.5 Powering Up and Checking LEDs

Perform the initial check of the array according to the following procedure:

1. Connect two AC (or DC) power cables to the power and fan modules on the back of the array.
2. Power on the array by turning on each power switch.

See [“Power-On Sequence” on page 4-23](#) for the power-up sequence to use when operating RAID arrays and expansion units. See [“Connecting the Sun StorEdge 3510 FC JBOD Array to a Host Computer” on page B-13](#) for the power-up sequence to use when operating standalone JBODs directly attached to hosts.

3. Check for the following LED activity:

All front panel LEDs turn solid green to indicate good operation.




FIGURE 4-5 Front Panel With LEDs Displayed

See the chapter [“Checking LEDs” on page 6-1](#) for more information about your array’s LEDs.

4.6 Reviewing and Changing SFP Ports (Optional)

Each I/O controller module has six ports that accept SFP transceivers. These ports are labeled FC0 through FC5. The default configurations do not include an SFP connector on every SFP port. You might want to add or rearrange SFPs, depending on the configuration mode (loop or point-to-point), number of planned host connections, the necessary number of redundant connections to hosts, and the number of expansion units needed.

The SFP is a single-port, optical SFP transceiver for multimode (short wave) or single-mode (long wave) use. They are compatible with the Small Form Factor Pluggable Multi-Sourcing Agreement (MSA, Sep. 2000), and 1x and 2x Fiber Channel. The optical connector used is the low-profile LC connector.

SFPs are Sun field-replaceable units (FRUs) and can be ordered from Sun Microsystems. These SFPs have been selected and tested to provide the necessary reliability and performance. SFPs from other vendors are not supported.

To see an example of an array that has SFP connectors in all the SFP ports of the RAID array and the upper expansion unit, refer to the illustration [“RAID Array Attached to Two Hosts and Two Expansion Units” on page 4-18](#).

To review various configuration options, see the discussions following [“Configuration Overview” on page 5-7](#). Refer also to the configuration options presented in the *Sun StorEdge 3000 Family Best Practices* for the Sun StorEdge 3510 FC Array.

4.6.1 Port Connectivity in a Dual-Controller Array

Drive channels connect to the internal drives in the array and can also connect to drives in external expansion units. Configured for redundancy, each controller in a dual-controller RAID array has two adjacent dedicated drive channels on one loop to load-balance I/O operations (see [FIGURE 4-6](#)). Each drive channel has two SFP ports that can be connected to expansion units. Drive channels 2 and 3 access all the disk drives and are interconnected to load-balance I/O operations.




FIGURE 4-6 Dedicated Drive Channels 2 (on the Upper Controller) and 3 (on the Lower Controller) in a Dual-Controller Array

The controller module in slot A (the upper slot) houses drive channel 2, which connects to the 12 internal disk drives through their A ports. The controller module in slot B (the lower slot) houses drive channel 3, which connects to the 12 internal disk drives through their B ports.

Host channels connect to host computers either directly or through storage switches or other devices. In a default dual-controller RAID configuration, each controller has four host channels, channels 0, 1, 4, and 5 (FIGURE 4-7). Port bypass circuits connect each pair of host SFP ports on a host channel; as a result, each host channel accesses both controllers.




FIGURE 4-7 Host Channels on a Dual-Controller Array

4.6.2 Default SFP Placement

In dual-controller arrays, SFPs are initially plugged into one of each pair of host and drive ports. The default port connections are:

- The upper I/O controller module has SFPs in the FC0, FC2, and FC4 ports.

- The lower I/O controller module has SFPs in the FC1, FC3, and FC5 ports.

This configuration provides connections to all four host channels as well as to both drive channels.




FIGURE 4-8 Default Dual-Controller SFP Placement

In the default single-controller array, SFPs are initially plugged into FC0, FC1, FC4, and FC5. No SFPs are plugged into the drive channels. This configuration is appropriate for connecting to up to four hosts or fibre switches, with no connection to expansion units.




FIGURE 4-9 Default Single-Controller SFP Placement

In a default JBOD expansion unit array, SFPs are initially plugged into the upper left port and lower right port.




FIGURE 4-10 Default JBOD Expansion Unit SFP Placement

4.6.3 Changing Your SFP Configuration

Fibre Channel arrays use SFP connectors to attach to hosts and expansion units. These SFP connectors resemble the one shown in [FIGURE 4-11](#), with a single connector at the end that plugs into an SFP port on the array or expansion unit chassis, and a duplex jack into which you insert a cable to make the connection.

- To make connections to an empty port, first slide the SFP connector into the port so that it connects firmly with the chassis. Then plug the fiber-optic cable's SFP connector into the duplex jack at the end of the SFP.
- To remove an SFP connector, make sure no cable is connected to it and then slide it out from the port.




FIGURE 4-11 Typical SFP Connector Used to Connect Cables to Chassis SFP Ports

4.7 Configuring a COM Port to Connect to a RAID Array

The RS-232 COM port on either controller module is used to configure and monitor the RAID array. It can be connected to a VT100 terminal or terminal emulation program, to a terminal server, or to the serial port of a server.

1. **Use a null modem serial cable to connect the COM port of the RAID array to the serial port on a host workstation.**

A null modem serial cable is included in your package contents.

2. **Set the serial port parameters on the workstation as follows:**

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

For platform-specific details, see the appendix that describes the operating system your server uses.

4.8 Setting an IP Address

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

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

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



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

There are two main reasons for placing your arrays on a private subnet:

- When your array is on a public network, it is susceptible to viruses, worms, and other malware attacks.
- A variety of security software is available to detect and mitigate these attacks. Some port-scanning and other security software can have an adverse impact on your ability to access data. In extreme cases, some of this software can cause Sun StorEdge 3000 family arrays to hang.

Since it is not possible to predict the side-effects of all current and future security software products, place your Sun StorEdge 3310 SCSI arrays and Sun StorEdge 3510 FC arrays on private subnets.

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

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

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

8. **Press Esc to continue.**

A confirmation prompt is displayed.

Change/Set IP Address ?

9. **Select Yes to continue.**

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

You are prompted to reset the controller.

10. Select Yes to reset the controller.

The controller takes a few minutes to format a small storage sector on each physical drive before logical drives can be successfully initialized.

4.9 Setting Up Out-of-Band Management Over Ethernet

The Ethernet connection enables you to configure and monitor RAID arrays and expansion units remotely by using the `telnet` command to access the firmware application on the array and by using the Sun StorEdge Configuration Service software.



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

1. To access the RAID array over an Ethernet connection, first set the IP address of the RAID array by using the RAID array's COM port and the RAID firmware. See [“Setting an IP Address” on page 4-14](#) for more information.
2. Connect the RAID array's Ethernet port on each controller to the network.

Note – In a dual-controller RAID array, be sure to connect both Ethernet ports to the network, which provides failover if one controller fails.

3. To use the firmware application program from the host server, connect to the IP address of the RAID array controller with the following command:

```
# telnet IP-address
```

4. Use the Solaris operating system `tip` command or a terminal emulation program to access the firmware application program. See [“Configuring a COM Port to Connect to a RAID Array” on page 4-14](#) for more information.

5. Press **Ctrl-L** to refresh the screen and view the Main Menu.

Note – If you reset the controller during a telnet session, you are disconnected from the RAID array. Use the `telnet` command to log back in to the array.

6. To connect the Sun StorEdge Configuration Service program (on a host server) to a RAID array that has an IP address, refer to the out-of-band management instructions in the *Sun StorEdge 3000 Family Configuration Service User's Guide*.

The same document's "Email and SNMP" appendix provides information about configuring Sun StorEdge Configuration Service software to use Simple Network Management Protocol (SNMP) traps and Management Information Bases (MIBs) to provide information to other out-of-band enterprise management software. The "Monitoring the Array" chapter explains the use of Sun StorEdge Configuration Service agents to redirect event messages into host system logs.

4.10 Cabling to Expansion Units

FIGURE 4-12 shows an example of a RAID array connected to two expansion units. You can connect up to eight expansion units to a RAID array in certain configurations. These configurations are described in the *Sun StorEdge 3000 Family Best Practices Manual*.

The connections to the expansion units are designed to keep all A drive ports on the same loop and all B drive ports on the same loop. Several cabling configurations are possible, each with its own advantages and disadvantages. Refer to the *Sun StorEdge 3000 Family Best Practices Manual* for the Sun StorEdge 3510 FC array for more information about the suitability of various configurations for various application requirements, as well as for information about configurations involving more than two expansion units.

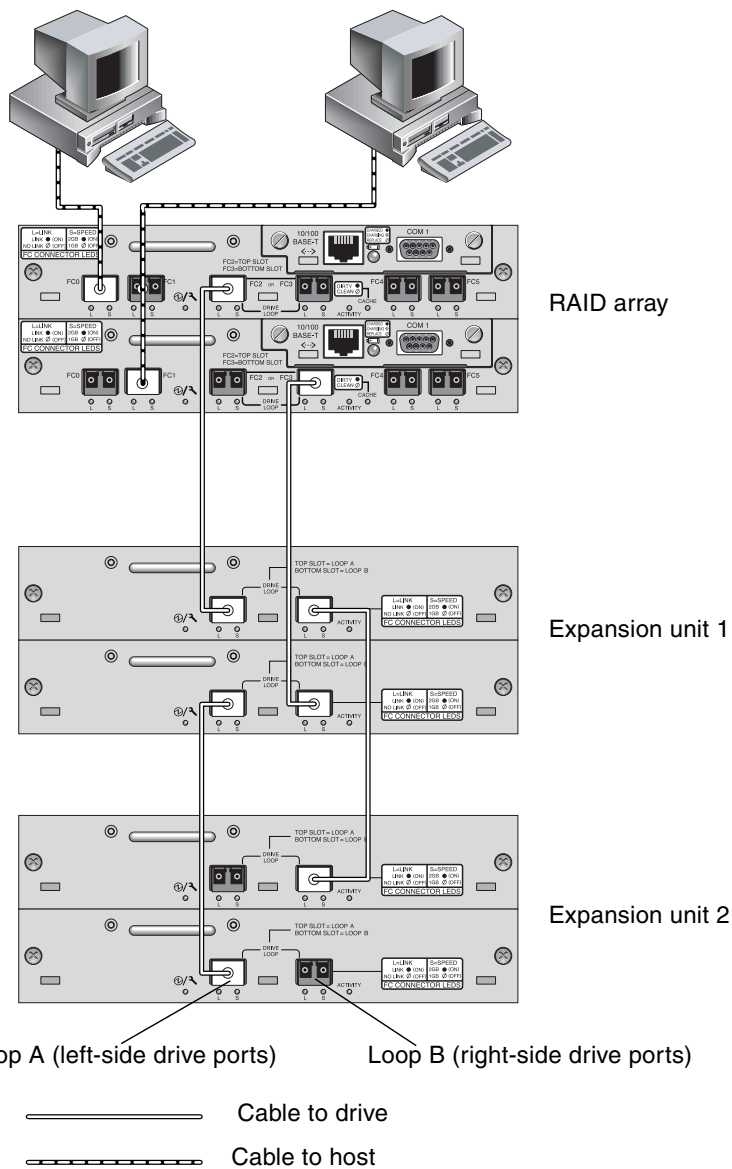


FIGURE 4-12 RAID Array Attached to Two Hosts and Two Expansion Units

In the RAID array, two of the unused SFP host ports could be used to provide redundant pathing to the two servers, and the other four unused SFP host ports could be connected to two more servers in a redundant configuration.

Similarly, you can connect expansion units to other channels (which are separate from channels 2 and 3) if you configure channels 0, 1, 4, or 5 as drive channels. For details, refer to [“Configuring FC Channels as Host or Drive \(Optional\)” on page 5-22.](#)



Caution – To avoid drive assignment conflicts, make sure that each connected array and expansion unit uses a different loop ID, as described in the following section.

4.10.1 Scaling Beyond 36 Disks

Note – High capacity Sun StorEdge 3510 FC array configurations with more than two expansion enclosures are supported, with certain limitations. For example, only the sequential cache optimization mode can be utilized in most situations. Other important limitations may exist.

Sun StorEdge 3510 FC arrays typically allow the connection of up to two expansion units to support a maximum of 36 disks. However, you can create larger configurations that support as many as eight expansion units and up to 108 disks if you use the guidelines in this section.

Carefully consider the limitations of these special high-capacity configurations. Using multiple Sun StorEdge 3510 FC arrays connected to the same SAN normally provides significantly better performance than one high-capacity configuration.

Limitations

- Random optimization should not be used with high-capacity configurations. Use sequential optimization instead. Random cache optimization greatly reduces the maximum number of supported disks.
- Only cabling configurations described in the *Sun StorEdge 3000 Family Best Practices Manual* for the Sun StorEdge 3510 FC array are supported when configuring high-capacity Sun StorEdge 3510 FC arrays.
- Maximize the size of each logical drive (up to 2048 Gbyte) before creating another logical drive. Otherwise the maximum number of eight logical drives might limit future expansion.
- Limiting the maximum number of expansion units to 7 provides much greater configuration flexibility. Doing so enables using up to 96 disks.

Refer to the *Sun StorEdge 3000 Family Best Practices Manual* for the Sun StorEdge 3510 FC array for cabling diagrams for configurations ranging from three to eight expansion units.

Note – Large configurations may require the use of one or more optional extended-length cables, part number X9732A. Other items may also be required. Refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for information about supported cables, SFPs, and other user-replaceable items.

4.11 Setting Loop IDs on Expansion Units

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

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

Setting the loop ID on an array requires dropping the front bezel out of the way and removing the small vertical plastic cap on the left side of the bezel that cover the left rackmount tab. Rackmount tabs are often referred to as “ears.”




FIGURE 4-13 Front Bezel and Front Bezel Locks of an Array

1. Use the provided key to unlock both bezel locks.
2. Grasp the front bezel cover on both sides and pull it forward and then down.
3. Remove the plastic cap from the left ear of the array.
 - a. Squeeze both sides of the cap at the top and the bottom.
 - b. Turn the cap toward the center of the array until it disengages and pull it free.



Caution – To avoid damage to the cap, do not pull the cap forward directly or pull from only its top or bottom.

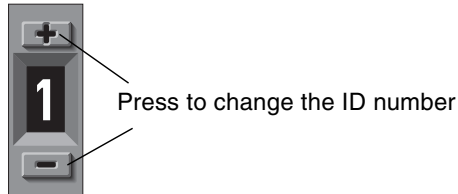


FIGURE 4-14 ID Switch Located on the Left Front Side of the Array or Expansion Unit

4. **Press the upper or lower switch button to change the ID number so that each connected RAID array and expansion unit uses a different loop ID.**

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



Caution – Ensure that the loop IDs of expansion units do not duplicate the loop IDs of other connected expansion units or RAID arrays.

The ID switch offers eight ID ranges. Each set contains 16 IDs (the last four IDs in each range are ignored). These ranges are shown in [TABLE 4-3](#).

TABLE 4-3 ID Switch Settings for Expansion Units

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

For an example of properly set loop IDs, consider the configuration shown in [FIGURE 4-12](#). You need to make sure the RAID array and each of the two expansion units have been assigned different loop IDs. If you set each of the loop ID switches so that the RAID array is assigned loop ID 0, expansion unit 1 has loop ID 1, and expansion unit 2 has loop ID 2. The range of IDs assigned to the drives is shown in [TABLE 4-4](#).

TABLE 4-4 Sample Array and Expansion Units With Different Loop IDs and Drive IDs

Fibre Channel Device	Loop ID Switch Setting	Range of Drive IDs
RAID array	0	0–15
Expansion unit 1	1	16–31
Expansion unit 2	2	32–47

5. Prepare the plastic left ear cap for replacement by aligning the inside round notches of the cap with the round cylindrical posts (ball studs) on the ear.
6. Push the top and bottom of the ear cap onto the ear, pressing in on the top side toward the center of the array first.
7. Continue pushing the top and bottom of the ear cap onto the ear, pressing on the side toward the outside of the array.
Do not use force when placing a cap on an ear.
8. Lift the bezel into position and press it onto the front of the chassis until it is flush with the front.

9. Use the key to lock both bezel locks.

4.12 Connecting Ports to Hosts

In a default array configuration, channels 0, 1, 4, and 5 are host channels, so you can directly connect an FC array to four host computers. SFP connectors are plugged into channels 0 and 4 on the upper controller and channels 1 and 5 on the lower controller for this purpose.

If you want to connect an FC array to more than four host computers without changing the default configuration, you can connect these four host channels to ports on storage switches in a storage area network (SAN) configuration.

Note – Connecting the Sun StorEdge 3510 RAID array to Fibre Channel HBAs on the same channel that use different speeds (1 Gbyte & 2 Gbyte) is not supported. You can, however, mix 1-Gbyte and 2-Gbyte Fibre Channel HBAs on different channels. This limitation is due to the design of the Sun StorEdge 3510 RAID array's port bypass circuitry and the inability of Fibre Channel to support auto-negotiation in a multi-drop loop configuration.

Use fiber-optic cables to connect from one to four host channels to Fibre Channel HBAs on your host computers or to other devices such as storage switches:

1. Connect a fiber-optic cable to an HBA or FC port on each host or storage switch you want to connect to the array.
2. Connect the SFP connector at the other end of each of these fiber-optic cables to host channel SFP connectors on the back of the array.

If there is no SFP connector in the port you want to use, first insert an SFP connector into the port as described in [“Changing Your SFP Configuration” on page 4-13](#).

4.13 Power-On Sequence

Power on the equipment in the following order, so that the host computer discovers all connected arrays:

- a. Expansion units
- b. FC arrays

c. Host computers

When you power on a controller, the Tip connection window displays a series of messages such as the following:

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

At this point if you type `skip` the channels will not be scanned. If you do not type anything, the channels are scanned and persistent group reservations are restored.

First-Time Configuration

This chapter summarizes the most common procedures used for first-time configuration.

This chapter covers the following topics:

- “Controller Defaults and Limitations” on page 5-2
 - “Planning for Reliability, Availability, and Serviceability” on page 5-2
 - “Dual-Controller Considerations” on page 5-3
 - “Single-Controller Considerations” on page 5-4
- “Battery Operation” on page 5-4
 - “Battery Status” on page 5-4
- “Software Management Tools” on page 5-6
- “Configuration Overview” on page 5-7
 - “Point-to-Point Configuration Guidelines” on page 5-9
 - “A SAN Point-to-Point Configuration Example” on page 5-11
 - “A DAS Loop Configuration Example” on page 5-15
- “Initial Configuration Steps” on page 5-18
 - “Viewing the Initial Firmware Windows” on page 5-19
 - “Checking Available Physical Drives” on page 5-20
 - “Configuring FC Channels as Host or Drive (Optional)” on page 5-22
 - “Choosing Loop or Point-to-Point Fibre Connection” on page 5-24
 - “Editing and Creating Additional Host IDs (Optional)” on page 5-25
 - “Selecting Sequential or Random Optimization” on page 5-27
 - “Reviewing Default Logical Drives and RAID Levels” on page 5-29
 - “Completing Basic Configuration” on page 5-31
 - “Creating Logical Drives (Optional)” on page 5-31
 - “Preparing for Logical Drives Larger Than 253 Gbyte” on page 5-38
 - “Changing a Logical Drive Controller Assignment (Optional)” on page 5-39
 - “Creating or Changing a Logical Drive Name (Optional)” on page 5-41
 - “Partitioning a Logical Drive (Optional)” on page 5-42
- “Mapping Logical Drive Partitions to Host LUNs” on page 5-45
 - “Setting Up Host Filter Entries” on page 5-52
- “Using Software to Monitor and Manage the Sun StorEdge 3510 FC Array” on page 5-60

Subsequent chapters in this manual describe further procedures used to complete the installation and configuration of FC arrays. The flexible architecture of the Sun StorEdge 3510 FC Array makes many configurations possible.

5.1 Controller Defaults and Limitations

This section describes default configurations and certain controller limitations.

5.1.1 Planning for Reliability, Availability, and Serviceability

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

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

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

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

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

5.1.2 Dual-Controller Considerations

The following controller functions describe the redundant controller operation.

- Both controllers must be exactly the same. They must operate with the same firmware version, the same size of memory, and the same number of host and drive channels. If one controller in a dual-controller configuration is replaced with a field-replaceable unit (FRU) controller, the array automatically compares the firmware versions of the two controllers. If they differ, the FRU controller firmware is automatically changed to that of the existing controller in the array.
- The controller firmware assumes that two controllers are available, or might be made available during operation at any time. In a 1U single-controller configuration, a 2U single-controller configuration, or a 2U dual-configuration, once the primary controller (which might be the only controller) is powered on, it begins to scan for a second controller. Until a second controller is discovered, which does not happen in a 1U single-controller configuration or a 2U single-controller configuration, the “Peripheral Device Status” for the redundant controller will show a Status of “Scanning.” This is correct behavior and enables the firmware to discover a second controller whenever it is added without the necessity of rebooting the primary controller.
- After booting in a redundant configuration, the controllers autonegotiate and designate one controller as primary and the other controller as secondary.
- The two controllers behave as one primary controller. Once the redundant configuration takes effect, user configurations and settings can be done only on the primary controller. The secondary controller then synchronizes with the configuration of the primary controller, making the configurations of the two controllers exactly the same.

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

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

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

5.1.3 Single-Controller Considerations

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

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



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

The Redundant Controller setting (“View and Edit Peripheral Devices → Set Peripheral Device Entry”) must remain enabled for single-controller configurations. This preserves the default primary controller assignment of the single controller. The controller status shows “scanning,” which indicates that the firmware is scanning for primary and secondary controller status and redundancy is enabled even though it is not used. There is no performance impact.

5.2 Battery Operation

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

5.2.1 Battery Status

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

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

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

TABLE 5-1 Battery Status Indicators

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

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

Note – The RAID controller has a temperature sensor which shuts off battery charging above 54 degrees C. When this happens, the battery status may be reported as BAD, but no alarm is written to the event log since no actual battery failure has occurred. This behavior is normal. As soon as the temperature returns to the normal range, battery charging resumes and the battery status is reported correctly. It is not necessary to replace or otherwise interfere with the battery in this situation.

For more information, see [“Environmental Requirements” on page 2-3](#) for the acceptable operating and nonoperating temperature ranges for your array.

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

5.2.2 Write-Back and Write-Through Cache Options

Unfinished writes are cached in memory in write-back mode. If power to the array is discontinued, data stored in the cache memory is not lost. Battery modules can support cache memory for 72 hours.

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

5.3 Software Management Tools

You can manage your array through an out-of-band or in-band connection.

5.3.1 Out-of-Band Connection

- With the out-of-band serial port connection, you can use a Solaris `tip` session or a Microsoft Windows terminal emulation program to access the firmware application. For details, see “[Setting Up the Serial Port Connection](#)” on page F-1.
- With the out-of-band Ethernet port connection, you can use the `telnet` command to access the firmware application. For details, see “[Setting an IP Address](#)” on page 4-14.



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

- With in-band host connections, you can use Sun StorEdge Configuration Service software or the command-line interface (CLI). Refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide* for in-band setup procedures and refer to the *Sun StorEdge 3000 3000 Family Software Installation Guide* for CLI installation and usage.

5.4 Configuration Overview

The Sun StorEdge 3510 FC Array is preconfigured and requires minimal configuration. All procedures can be performed by using the COM port. You can perform all procedures except the assignment of an IP address through an Ethernet port connection to a management console.

The typical sequence of steps for completing a first-time configuration of the array is to:

1. Make sure that the mounting of the array on a rack, cabinet, desk, or table is complete.
2. Set up the serial port connection. See [“Configuring a COM Port to Connect to a RAID Array” on page 4-14.](#)
3. Set up an IP address for the controller. See [“Setting an IP Address” on page 4-14.](#)
4. Check available physical drives. [“Checking Available Physical Drives” on page 5-20.](#)
5. Determine whether sequential or random optimization is more appropriate for your applications and configure your array accordingly. See [“Selecting Sequential or Random Optimization” on page 5-27.](#)
6. (Optional) Configure host channels as drive channels. See [“Configuring FC Channels as Host or Drive \(Optional\)” on page 5-22.](#)
7. Confirm or change the Fibre Connection Option (point-to-point or loop). See [“Choosing Loop or Point-to-Point Fibre Connection” on page 5-24.](#)
8. Revise or add host IDs on host channels. [“Editing and Creating Additional Host IDs \(Optional\)” on page 5-25.](#)

The IDs assigned to controllers take effect only after the controller is reset.

9. (Optional) Remove default logical drives and create new logical drives. See [“Creating Logical Drives \(Optional\)” on page 5-31.](#)
10. (Optional) In dual-controller configurations only, assign logical drives to the secondary controller to load-balance the two controllers. [“Changing a Logical Drive Controller Assignment \(Optional\)” on page 5-39.](#)



Caution – In single-controller configurations, do not disable the Redundant Controller Function and do not set the controller as a secondary controller. The primary controller controls all firmware operations and must be the assignment

of the single controller. If you disable the Redundant Controller Function and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module becomes inoperable and will need to be replaced.

Note – While the ability to create and manage logical volumes remains a feature of Sun StorEdge 3000 family FC and SCSI RAID arrays for legacy reasons, the size and performance of physical and logical drives have made the use of logical volumes obsolete. Logical volumes are unsuited to some modern configurations, such as Sun Cluster environments, and do not work in those configurations. Avoid using them and use logical drives instead. For more information about logical drives, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide*.

11. **(Optional) Partition the logical drives.** [“Partitioning a Logical Drive \(Optional\)” on page 5-42.](#)
 12. **Map each logical drive partition to an ID on a host channel, or apply a host LUN filter to the logical drives.** See [“First Steps in Mapping a Partition to a LUN” on page 5-48](#) for more information.
-

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

For information about different operating environment procedures, refer to:

- [Appendix F, “Configuring a Sun Server Running the Solaris Operating Environment” on page F-1](#)
- [Appendix G, “Configuring a Windows 2000 Server or Windows 2000 Advanced Server” on page G-1](#)
- [Appendix H, “Configuring a Linux Server” on page H-1](#)
- [Appendix I, “Configuring an IBM Server Running the AIX Operating Environment” on page I-1](#)
- [Appendix J, “Configuring an HP Server Running the HP-UX Operating Environment” on page J-1](#)
- [Appendix K, “Configuring a Windows NT Server” on page K-1](#)

13. **Reset the controller.**

Configuration is complete.

Note – Resetting the controller can result in occasional host-side error messages such as parity error and synchronous error messages. No action is required and the condition corrects itself as soon as reinitialization of the controller is complete.

14. **Save the configuration to a disk.** See [“Saving Configuration \(NVRAM\) to a Disk” on page 5-59.](#)
15. **Make sure that the cabling from the RAID array to the hosts is complete.**

Note – You can reset the controller after each step or at the end of the configuration process.



Caution – Avoid using in-band and out-of-band connections at the same time to manage the array. Otherwise, conflicts between multiple operations can cause unexpected results.

5.4.1 Point-to-Point Configuration Guidelines

Remember the following guidelines when implementing point-to-point configurations in your array and connecting to Fabric switches:

- The default mode is “Loop only.” You must change the Fibre Channel Connection mode to “Point-to-point only” with the firmware application.



Caution – If you keep the default loop mode and connect to a Fabric switch, the array automatically shifts to public loop mode. As a result, communication between the array and the switched Fabric runs in half duplex (send or receive) instead of providing the full duplex (send and receive) performance of point-to-point mode.

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

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

However, if you ignore this warning and add an ID to the other controller, the RAID controller does not allow a login as a Fabric-Loop (FL) port since this would be illegal in a point-to-point configuration.

- The firmware application allows you to add up to eight IDs per channel (four IDs on each controller), which forces the Fabric switch port type to become Fabric-Loop. To ensure F-port behavior (full Fabric/full duplex) when attaching to a switch, only one ID must be present on each channel and the array port must be set to point-to-point mode.
- Do not connect more than one port per channel on an array to a Fabric switch.



Caution – In point-to-point mode or in public loop mode, only one switch port is allowed per channel. Connecting more than one port per channel to a switch can violate the point-to-point topology of the channel, force two switch ports to “fight” over an AL_PA (arbitrated loop physical address) value of 0 (which is reserved for loop to Fabric attachment), or both.

- With four host channels and four host IDs, you should load-balance the host ID setup so that half the IDs are on the primary controller and half the IDs are on the secondary controller. When setting up LUNs, map each LUN to either two PIDs or two SIDs. The hosts are in turn dual-pathed to the same two switched Fabrics. When attaching the cables, for a LUN-mapped channel pair, make sure that the first channel is connected to the upper port and the second channel is connected to the lower port of the second channel.

For example, to provide redundancy, map half of the LUNs across Channel 0 (PID 40) and Channel 4 (PID 42), and then the other half of your LUNs across Channel 1 (SID 41) and Channel 5 (SID 43)

- Point-to-point mode has a maximum of 128 LUNs per array. In a redundant configuration, 32 LUNs are dual-mapped across two channels on the primary controller, and another 32 LUNs are dual-mapped across the secondary controller, for a total of 64 distinct LUNs.
- To use more than 64 LUNs, you must change to “Loop only” mode, add host IDs to one or more channels, and add 32 LUNs for each additional host ID.

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

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

5.4.2 A SAN Point-to-Point Configuration Example

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

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

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

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

Important rules to remember are:

- A single logical drive can be mapped to only one controller, either the primary controller or the secondary controller.
- In a point-to-point configuration, only one host ID per channel is allowed. The host ID can be assigned to the primary controller and be a PID, or it can be assigned to the secondary controller and be a SID.
- If you have two switches and set up multipathing (to keep all logical drive connections operational for any switch failure or the removal of any I/O controller module), ensure that each logical drive is mapped to two ports, one on each I/O controller module, and on two channels. The cables from the two ports mapped to each logical drive must be cabled to two separate switches. See [FIGURE 5-1](#) for an example of this configuration.


[FIGURE 5-1](#) shows the channel numbers (0, 1, 4, and 5) of each host port and the host ID for each channel. N/A means that the port does not have a second ID assignment. The primary controller is in the top I/O controller module, and the secondary controller is in the bottom I/O controller module.

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

- The port bypass circuit on each channel connects the upper and lower ports on the same channel and provides access to both controllers at the same time.
- If there are two host connections to the upper and lower ports on Channel 0, and one host connection is removed, the other host connection remains operational.
- Therefore, if you have a redundant multipathing configuration in which you have two host connections to each logical drive and one connection fails, the remaining path maintains a connection to the logical drive.

In [FIGURE 5-1](#), with multipathing software to reroute the data paths, each logical drive remains fully operational when the following conditions occur:

- One switch fails or is disconnected, and the logical drive is routed to the second switch. For example, if switch 0 fails, switch 1 automatically accesses logical drive 0 through the cabling to the lower port on PID 41.
- One I/O controller module fails, and all the host IDs for that controller are reassigned (moved) to the second I/O controller module. For example, if the upper I/O controller module is removed, host IDs 40 and 41 are automatically moved to the lower module and are managed by the second controller.
- An I/O controller module fails or one cable is removed from an I/O controller module, and all I/O traffic to the disconnected channel is rerouted through the second port/host LUN assigned to the logical drive. For example, if you remove the cable to channel 4, the data path for logical drive 1 switches to the port on channel 5.



Map LG0 to PIDs 40 and 41

Map LG1 to SIDs 50 and 51




FIGURE 5-1 A Point-to-Point Configuration With a Dual-Controller Array and Two Switches

Note – This illustration shows the default controller locations; however, the primary controller and secondary controller locations can occur on either slot and depend on controller resets and controller replacement operations.

TABLE 5-2 summarizes the primary and secondary host IDs assigned to logical drives 0 and 1, based on [FIGURE 5-1](#).

TABLE 5-2 Example Point-to-Point Configuration With Two Logical Drives in a Dual-Controller Array

Task	Logical Drive	LUN IDs	Channel Number	Primary ID Number	Secondary ID Number
Map 32 partitions of LG0 to CH0	LG 0	0-31	0	40	N/A
Duplicate-map 32 partitions of LG0 to CH1	LG 0	0-31	1	41	N/A
Map 32 partitions of LG1 to CH4	LG 1	0-31	4	N/A	50
Duplicate-map 32 partitions of LG1 to CH5	LG 1	0-31	5	N/A	51

Perform the following steps, which are described in more detail later in this manual, to set up a typical point-to-point SAN configuration based on [FIGURE 5-1](#).

- 1. Check the position of installed SFPs. Move them as necessary to support the connections needed.**
- 2. Connect expansion units, if needed.**
- 3. Create at least two logical drives (logical drive 0 and logical drive 1) and configure spare disks.**
For half of the logical drives, keep the logical drives assigned to the primary controller by default. For the other half of the logical drives, assign the logical drives to the secondary controller to load-balance the I/O.
- 4. Create up to 32 partitions (LUNs) in each logical drive, for each server.**
- 5. Change the Fibre Connection Option to “Point to point only.”**
- 6. For ease of use in configuring LUNs, change the host IDs on the four channels to be the following assignments:**
Channel 0: PID 40 (assigned to the primary controller)
Channel 1: PID 41 (assigned to the primary controller)
Channel 4: SID 50 (assigned to the secondary controller)
Channel 5: SID 51 (assigned to the secondary controller)



Caution – Do not use the command, “Point to point preferred, otherwise loop.” This command is reserved for special use and should be used only if directed by technical support.

- 7. Map logical drive 0 to channels 0 and 1 of the primary controller.**
Map LUN numbers 0 through 31 to the single ID on each host channel.

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

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

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

9. Connect the first switch to ports 0 and 4 of the upper controller.

10. Connect the second switch to ports 1 and 5 of the lower controller.

11. Connect each server to each switch.

12. Install and enable multipathing software on each connected server.

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

5.4.3 A DAS Loop Configuration Example

The typical direct attached storage (DAS) configuration shown in [FIGURE 5-2](#) includes four servers, a dual-controller array, and two expansion units. Expansion units are optional.

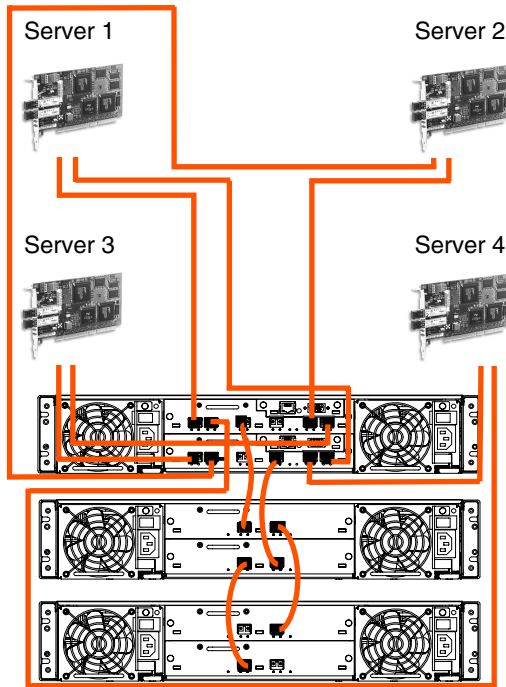


FIGURE 5-2 A DAS Configuration With Four Servers, a Dual-Controller Array, and Two Expansion Units

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

1. Establish two connections between a server and a Sun StorEdge 3510 FC Array,
2. Install and enable multipathing software on the server,
3. Map the logical drive the server is using to the controller channels the server is connected to.

DAS configurations are typically implemented using a Fabric loop (FL_port) mode. A loop configuration example is described under [“A DAS Loop Configuration Example”](#) on page 5-15.

Fabric loop (FL_port) connections between a Sun StorEdge 3510 FC Array and multiple servers allow up to 1024 LUNs to be presented to servers.

For guidelines on how to create 1024 LUNs, see [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)”](#) on page 5-47.

Perform the following steps, which are described in more detail later in this manual, to set up a DAS loop configuration based on [FIGURE 5-2](#).

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

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

2. Connect expansion units, if needed.

3. Create at least one logical drive per server, and configure spare disks as needed.

4. Create one or more logical drive partitions for each server.

5. Confirm that the Fibre Connection Option is set to “Loop only.”

Do not use the “Loop preferred, otherwise, point to point” option, which should not be used for this product.



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

6. Set up to eight IDs on each channel, if needed.

TABLE 5-3 Example Primary and Secondary ID Numbers in a Loop Configuration With Two IDs per Channel

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

7. Map logical drive 0 to channels 0 and 5 of the primary controller.

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

9. Map logical drive 2 to channels 0 and 5 of the primary controller.

10. Map logical drive 3 to channels 1 and 4 of the secondary controller.

11. Connect the first server to port 0 of the upper controller and port 5 of the lower controller.

12. Connect the second server to port 1 of the lower controller and port 4 of the upper controller.
13. Connect the third server to port 0 of the lower controller and port 5 of the upper controller.
14. Connect the fourth server to port 1 of the upper controller and port 4 of the lower controller.
15. Install and enable multipathing software on each connected server.

5.4.4 Larger Configurations

Larger configurations are possible using additional expansion units. Up to eight expansion units are supported when connected to a Sun StorEdge 3510 FC RAID array. Several such configurations are possible. For more detailed information, and for suggestions about the most appropriate configurations for your applications and environment, refer to the *Sun StorEdge 3000 Family Best Practices Manual* for the Sun StorEdge 3510 FC array.

5.5 Initial Configuration Steps

The topics in this section present both required and commonly used optional procedures, which apply to both the point-to-point and loop mode configurations in most cases.

Note – If you want to create logical volumes, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide*. Logical volumes are not widely used since they do not allow partitions and limit the number of available LUNs.

Most of the configuration you do involves using firmware menus to change settings on the array. However, each host platform involves some initial configuration as well. Refer to the appendix for your host platform for instructions on how to connect your host to an array, for host-specific instructions on recognizing and formatting LUNs including modifying host configuration files, and for other platform-specific details.

5.5.1 Viewing the Initial Firmware Windows

You see the initial controller screen when you first access the RAID controller firmware.

The initial screen is displayed when the RAID controller is powered on.

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

1. Use the up and down arrow keys to select the VT100 terminal emulation mode, and then press Return to enter the Main Menu.

See [“Introducing Key Screens and Commands” on page 7-2](#) for detailed information about understanding and using the initial firmware screen.

2. Use the following keys to navigate within the application:

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

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

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

- Proceed to configure the array using options from the Main Menu as described in the rest of this chapter.

```

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

```

FIGURE 5-3 Firmware Main Menu

5.5.2 Checking Available Physical Drives

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

- Choose “view and edit scsi Drives” from the Main Menu.

A list of all the installed physical drives is displayed.

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

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

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

When the power is on, the controller scans all hard drives that are connected through the drive channels. If a hard drive was connected after the controller completed initialization, use the “Scan scsi drive” menu option to let the controller recognize the newly added hard drive and configure it.

Caution – Using the “Scan scsi drive” menu option to recognize and configure a drive removes its assignment to any logical drive. All data on that drive is lost.

- To review more information about a drive, highlight it and press Return. Then choose “View drive information” and press Return to view details about that drive.

Quick view view view view view view view view view view view	Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID	
	2<3>	0	0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
	View drive information						0	ON-LINE	SEAGATE ST336752FSUN36G
	Scan scsi drive						0	ON-LINE	SEAGATE ST336752FSUN36G
	set slot Number						0	ON-LINE	SEAGATE ST336752FSUN36G
	add drive Entry						0	ON-LINE	SEAGATE ST336752FSUN36G
	Identify scsi drive						0	ON-LINE	SEAGATE ST336752FSUN36G
	clone Failing drive						0	ON-LINE	SEAGATE ST336752FSUN36G
	disk Reserved space - 256 mb						0	ON-LINE	SEAGATE ST336752FSUN36G
	2<3>	5	5	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
	2<3>	6	6	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	
	2<3>	7	7	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	

Additional information is displayed about the drive you selected.

Quick view view view view view view view view view view view	Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID	
	2<3>	0	0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
	View drive information						0	ON-LINE	SEAGATE ST336752FSUN36G
	Revision Number						0205	36752FSUN36G	
	Serial Number						3E10MNP900007251	36752FSUN36G	
	Disk Capacity (blocks)						71132958	36752FSUN36G	
	Node Name(WWN)						20 00 00 04 CF A7 CE E4	36752FSUN36G	
	Redundant Loop ID						0	36752FSUN36G	
	2<3>	5	5	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
	2<3>	6	6	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	
	2<3>	7	7	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	

5.5.3 Configuring FC Channels as Host or Drive (Optional)

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

- Default channel settings are:
 - CH 0, CH 1, CH 4, and CH 5 = Host channels
 - CH 2 and CH 3 = Drive channels
- Channels 2 and 3 (CH 2 and 3) must be drive channels.
- Channels 0, 1, 4, and 5 can be drive or host channels. The most common reason to change a host channel to a drive channel is to attach an expansion unit to a RAID array.

To change a host channel to a drive channel, reconfigure the channel according to the following procedure:

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

Channel information is displayed.

```

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

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Serial
1	Host	NA	42	2 GHz	Serial	F	NA	2 GHz	Serial
2<3;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
3<2;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
4	Host	44	NA	2 GHz	Serial	F	NA	2 GHz	Serial
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Serial

Note – The Mode column for at least one channel must include the RCC abbreviation for Redundant Controller Communications.

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

3. Choose "channel Mode."

A menu of channel mode options is displayed.

< Main Menu > Quick installation view and edit Logical drives view and edit logical Volumes										
Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid	
0	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Serial	
1	channel Mode view and edit scsi Id view chip inFormation view channel host-id/Wwn View device port name list<wwpn> Data rate					1	F	NA	2 GHz	Serial
2<						1	F	NA	2 GHz	Serial
3<						1	F	NA	2 GHz	Serial
4						1	F	NA	2 GHz	Serial
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Serial	

4. Modify the channel to suit your requirements.

Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for more information about modifying channels.

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

A confirmation message is displayed:

NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally. Do you want to reset the controller now?

6. Choose Yes to reset the controller.

5.5.4 Choosing Loop or Point-to-Point Fibre Connection

To confirm or change the fibre connection for the array, perform the following steps:

1. Choose “view and edit Configuration parameters → Host-side SCSI Parameters → Fibre Connection Option.”

```
< Main Menu >
Quick installation
view
view Maximum Queued I/O Count - 256
view LUNs per Host SCSI ID - 32
view Max Number of Concurrent Host-LUN Connection - 1024
view Number of Tags Reserved for each Host-LUN Connection - Def(32)
view Peripheral Device Type Parameters
view Host Cylinder/Head/Sector Mapping Configuration
v C Fibre Connection Option - Loop only
v C
v C
v C
v C
Host-side SCSI Parameters
Drive-side SCSI Parameters
Disk Array Parameters
Redundant Controller Parameters
Controller Parameters
DMEP Parameters
```

2. If you want to view or change the Fibre Connection Option, choose either “Loop only” or “Point to point only” and press Return.

A confirmation dialog is displayed.

3. Choose Yes to confirm.

```
< Main Menu >
Quick installation
view
view Maximum Queued I/O Count - 256
view LUNs per Host SCSI ID - 32
view Max Number of Concurrent Host-LUN Connection - 1024
view Number of Tags Reserved for each Host-LUN Connection - Def(32)
view Peripheral Device Type Parameters
view Host Cylinder/Head/Sector Mapping Configuration
v C Fibre Connection Option - Loop only
v C
v C
v C
v C
Host Loop only
Drive Point to point only
```



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

A confirmation message is displayed:

NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally. Do you want to reset the controller now?

4. Choose Yes to reset the controller.

5.5.5 Editing and Creating Additional Host IDs (Optional)

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

Default host channel IDs are displayed in [TABLE 5-4](#).

TABLE 5-4 Default Host Channel IDs

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

The number of host IDs depends on the configuration mode:

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

Typically host IDs are distributed between the primary and secondary controllers to load-balance the I/O in the most effective manner for the network.

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

- Edit each host ID number to change the target number of each controller host channel that is seen by the host
- Add additional host ID numbers for loop configurations

Note – To map 1024 partitions in loop mode, you must add additional host IDs so that 32 IDs are mapped to the array’s channels. Several configurations are possible, such as eight IDs mapped to each of the four host channels or sixteen IDs mapped to two channels and none to the other two. For more information, see [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)”](#) on page 5-47.

To add a unique host ID number to a host channel, perform the following steps.

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

```

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

Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Serial
1	Host	NA	42	2 GHz	Serial	F	NA	2 GHz	Serial
2<3;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
3<2;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
4	Host	NA	44	2 GHz	Serial	F	NA	2 GHz	Serial
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Serial

2. Select the host channel whose Primary/Secondary ID you want to edit and press Return.
3. Choose “view and edit scsi Id.”
A dialog is displayed that says “No SCSI ID Assignment - Add Channel SCSI ID?”
4. Choose Yes to confirm.
“Primary Controller” and “Secondary Controller” are displayed in a menu.
5. Select the primary or secondary controller to which you want to add a host ID.
By default, channel 0 has a primary ID (PID) and no secondary ID (SID), while channel 1 has a SID but no PID.
6. Select an ID number for that controller and press Return.

7. Confirm your selection by choosing Yes and pressing Return.

Main Menu										
Quick installation										
view and edit Logical drives										
view and edit logical Volumes										
Ch1	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	Cur	ID 32
0	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Ser	ID 33
1	ID 40 <Primary Controller>				Serial	F	NA	2 GHz	Ser	ID 34
2<3;D	Add Channel SCSI ID				Serial	F	NA	2 GHz	Ser	ID 35
3<2;D	Primary Controller				Serial		Add Primary Controller SCSI ID			ID 36
4	Secondary Controller				Serial		Yes	No		ID 37
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Ser	ID 38
										ID 39
										ID 40
										ID 41
										ID 42
										ID 43
										ID 44
										ID 45
										ID 46
										ID 47
										ID 48
										ID 49
										ID 50

A confirmation message is displayed:

NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally. Do you want to reset the controller now?

8. Choose Yes to Reset the controller.

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

5.5.6 Selecting Sequential or Random Optimization

Before creating or modifying logical drives, select the optimization mode for all logical drives you create. The optimization mode determines the block size used when writing data to the drives in an array.

- Sequential I/O uses large blocks of data:
 - 128 Kbyte for RAID 0, 1, and 5 logical drives
 - 16 Kbyte for RAID 3 logical drives)
- Random I/O uses small blocks of data:
 - 32 Kbyte) for RAID 0, 1, and 5 logical drives
 - 4 Kbyte for RAID 3 logical drives

Note – Your array is preconfigured for Sequential Optimization. If Random Optimization is most appropriate for your use, you will need to delete all of the preconfigured logical drives, change the optimization mode, and then create new logical drives.

The type of application the array is working with determines whether random or sequential I/O should be applied. Video and imaging application’s I/O size can be 128, 256, 512 Kbyte, or up to 1 Mbyte, so the application reads and writes data to and

from the drive as large-block, sequential files. Database and transaction-processing applications read and write data from the drive as small-block, randomly accessed files.

There are two limitations that apply to the optimization modes:

- One optimization mode must be applied to all logical drives in an array.
- Once the optimization mode is selected and data written to logical drives, the only way to change the optimization mode is to back up all data to another location, delete all logical configurations of drives, reconfigure the logical drive with the new optimization mode, and reboot the array.

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

For more information about optimization modes, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array.

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

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

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

TABLE 5-5 Maximum Number of Disks per Logical Drive for a 2U Array

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

TABLE 5-6 Maximum Usable Capacity (Gbyte) per Logical Drive for a 2U Array

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

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

5.5.7 Reviewing Default Logical Drives and RAID Levels

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

A drive can be assigned as the local spare drive to one specified logical drive, or as a global spare drive that is available to all logical drives on the RAID array.

Spares can be part of an automatic array rebuild.

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

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

- A logical drive can be further divided into a maximum of 128 partitions.
- In loop mode, the maximum number of partitions is 1024 partitions per RAID array. To create a maximum of 1024 partitions, see [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)”](#) on page 5-47.
- In point-to-point mode, the maximum number of partitions is 64 partitions in a redundant configuration.

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

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

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

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

The following table highlights the RAID levels available.

TABLE 5-7 RAID Level Definitions

RAID Level	Description
RAID 0	Striping without data redundancy; provides maximum performance.
RAID 1	Mirrored or duplexed disks; for each disk in the system, a duplicate disk is maintained for data redundancy. It requires 50% of total disk capacity for overhead.
RAID 3	Striping with dedicated parity. Parity is dedicated to one drive. Data is divided into blocks and striped across the remaining drives.
RAID 5	Striping with distributed parity; this is the best-suited RAID level for multitasking or transaction processing. The data and parity are striped across each drive in the logical drive, so that each drive contains a combination of data and parity blocks.
NRAID	NRAID is an older configuration that is rarely used and is not recommended.
RAID 1+0	RAID 1+0 combines RAID 1 and RAID 0—mirroring and disk striping. RAID 1+0 allows multiple drive failures because of the full redundancy of the hard disk drives. If four or more hard disk drives are chosen for a RAID 1 logical drive, RAID 1+0 is performed automatically.
RAID (3+0)	A logical volume with several RAID 3 member logical drives.
RAID (5+0)	A logical volume with several RAID 5 member logical drives.

Note – Logical volumes are unsuited to some modern configurations, such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information see [“Configuration Overview” on page 5-7](#).

For more information about logical drives, spares, and RAID levels, refer to the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for your array.

5.5.8 Completing Basic Configuration

In a point-to-point configuration, the last required step is mapping the logical drives to host LUNs.

In loop mode, you have additional options to pursue, if needed, in addition to the mapping requirement:

- Optionally, define any additional partitions for each logical drive. See [“Partitioning a Logical Drive \(Optional\)” on page 5-42](#).
- Optionally, add host FC IDs and more logical drives to create 1024 LUNs. See:
 - [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)” on page 5-47](#)
 - [“Editing and Creating Additional Host IDs \(Optional\)” on page 5-25](#)
 - [“Creating Logical Drives \(Optional\)” on page 5-31](#)

For the procedure on the required mapping to LUNs, see [“First Steps in Mapping a Partition to a LUN” on page 5-48](#).

Note – Alternatively, you can use the graphical user interface described in the *Sun StorEdge 3000 Family Configuration Service User’s Guide* to map the logical drives to host LUNs.

5.5.9 Creating Logical Drives (Optional)

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

If you prefer a different configuration, use the procedure described in this section to modify the RAID level or to add more logical drives. In this procedure, you configure a logical drive to contain one or more hard drives based on the desired RAID level, and partition the logical drive into additional partitions.

Note – Logical volumes are unsuited to some modern configurations, such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information see [“Configuration Overview” on page 5-7](#).

Note – If you want to create 1024 LUNs in loop mode, you need eight logical drives with each having 128 partitions.

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

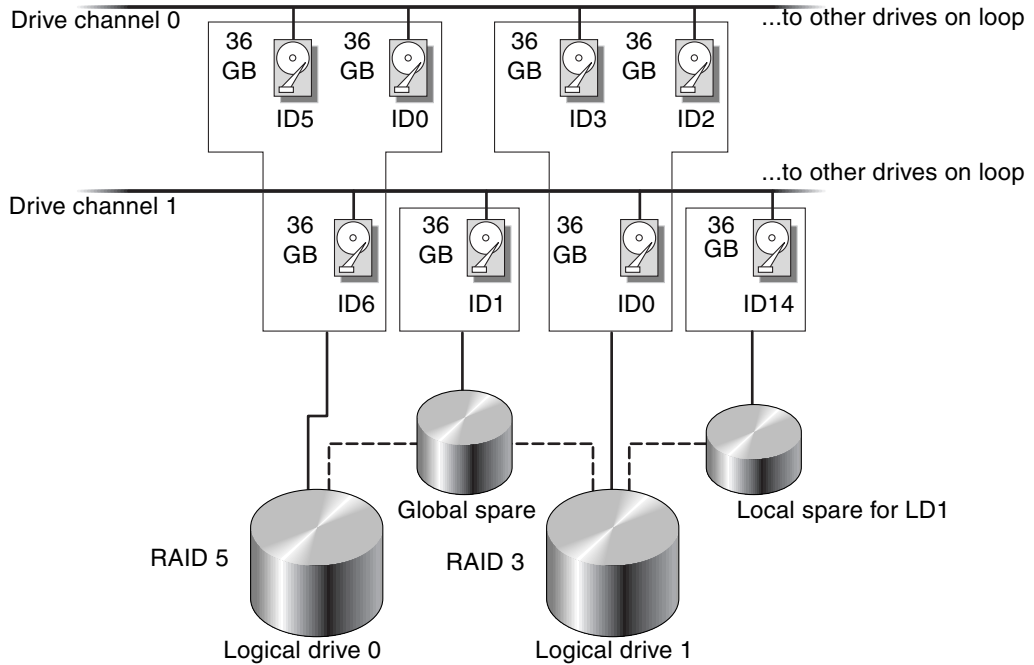


FIGURE 5-4 Example of an Allocation of Local and Spare Drives in Logical Configurations

Note – To reassign drives and add additional local or global spares on your preconfigured array, you must first unmap and then delete the existing logical drives, and then create new logical drives. For more information on deleting a logical drive, see [“Deleting Logical Drives”](#) on page 7-6.

Create a logical drive with the following steps:

1. Choose **“view and edit Logical drives”** from the Main Menu.

```

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


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

Q	LG	Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
U	0	*	2<3>	112	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	1		2<3>	113	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	2		2<3>	114	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	3		2<3>	119	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	4		2<3>	120	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	5		2<3>	121	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	6		2<3>	122	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	7		2<3>	123	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G

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

To deselect a drive, press Return again on the selected drive. The asterisk disappears.

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

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

Q	LG	Slot	Ch1	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID
U	0	*	2<3>	112	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	1	*	2<3>	113	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	2	*	2<3>	114	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	3		2<3>	119	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	4		2<3>	120	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	5		2<3>	121	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	6		2<3>	122	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G
U	7		2<3>	123	34732	200MB	NONE	FRMT DRU	SEAGATE ST336752FSUN36G

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

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

4			NONE							
	Maximum Drive Capacity : 34476MB Assign Spare Drives Disk Reserved Space: 256 MB Logical Drive Assignments									

4			NONE							
	Slot	Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID		
		0	12	34732	160MB	NONE	FRMT DRV	SEAGATE ST336605L SUN36G		
		0	13	34732	160MB	NONE	NEW DRV	SEAGATE ST336605L SUN36G		

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

7. (Optional) For dual-controller configurations only, choose “Logical Drive Assignments” to assign this logical drive to the secondary controller.

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



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

Q	LG	ID	LU	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
U	0			NONE							
U	Maximum Drive Capacity : 2000MB Assign Spare Drives Disk Reserved Space: 256 MB Logical Drive Assignments										
U	Redundant Controller Logical Drive Assign to Secondary Controller ? Yes No										
U	4										
U	5										
S	6			NONE							
S	7			NONE							

If you use two controllers for the redundant configuration, a logical drive can be assigned to either of the controllers to balance the workload. Logical drive assignments can be changed later, but require a controller reset to take effect.

8. Choose “view and edit Logical drives” to view the first created logical drive (P0) on the first line of the status window.

Q	LG	ID	LU	RAID	Size<MB>	Status	O	#LN	#SB	#FL	NAME
U	P0	4DB84961	NA	RAID5	4000	GOOD	S	3	0	0	
U	1			NONE							
U	2			NONE							
U	3			NONE							
U	4			NONE							
U	5			NONE							
U	6			NONE							
U	7			NONE							

5.5.10 Preparing for Logical Drives Larger Than 253 Gbyte

The Solaris operating environment requires drive geometry for various operations, including `newfs`. For the appropriate drive geometry to be presented to the Solaris operating environment for logical drives larger than 253 Gbyte, change the default settings to “< 65536 Cylinders” and “255 Heads” to cover all logical drives over 253 GB and under the maximum limit. The controller automatically adjusts the sector count, and then the operating environment can read the correct drive capacity.

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

TABLE 5-8 Cylinder and Head Mapping for the Solaris Operating Environment

Logical Drive Capacity	Cylinder	Head	Sector
< 253 GB	variable (default)	variable (default)	variable (default)
253 GB - 1 TB	< 65536 Cylinders	255	variable (default)

Note – Earlier versions of the Solaris operating environment do not support drive capacities larger than 1 terabyte.

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


TABLE 5-9 Cylinder and Head Mapping for the Solaris Operating Environment

Logical Drive Capacity	Cylinder	Head	Sector
< 253 GB	variable (default)	variable (default)	variable (default)
253 GB - 1 TB	< 65536 Cylinders *	255 *	variable (default)

* These settings are also valid for all logical drives under 253 GBytes. After settings are changed, they apply to all logical drives in the chassis.

To revise the Cylinder and Head settings, perform the following steps.

1. Choose “view and edit Configuration parameters → Host-Side SCSI Parameters → Host Cylinder/Head/Sector Mapping Configuration → Sector Ranges - Variable → 255 Sectors → Head Ranges - Variable.”
2. Specify “255 Heads.”
3. Choose “Cylinder Ranges - Variable → < 65536 Cylinders.”



Refer to *Sun StorEdge 3000 Family RAID Firmware User's Guide* for more information about firmware commands used with logical drives.

5.5.11 Changing a Logical Drive Controller Assignment (Optional)

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

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



Caution – In single-controller configurations, do not disable the Redundant Controller Function and do not set the controller as a secondary controller. The primary controller controls all firmware operations and must be the assignment of the single controller. If you disable the Redundant Controller Function and reconfigure the controller with the Autoconfigure option or as a secondary controller, the controller module becomes inoperable and will need to be replaced in a single-controller configuration.

After a logical drive has been created, it can be assigned to the secondary controller. Then the host computer associated with the logical drive can be mapped to the secondary controller (see [“First Steps in Mapping a Partition to a LUN” on page 5-48](#)).

To change a logical drive controller assignment:

1. Choose “view and edit Logical drives” from Main Menu.

```

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

2. Select the drive you want to reassign and press Return.
3. Choose “logical drive Assignments” and press Return.

Q	LG	ID	LW	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S	View scsi drives				GOOD	S	6	-	0	
U	P	Delete logical drive				GOOD	S	6	-	0	
U	P	Partition logical drive									
U	S	logical drive Name									
U	S	logical drive Assignments				GOOD	S	6	-	0	
U		Expand logical drive									
U		add Scsi drives									
U		copy and replace drive									
	6			NONE							
	7			NONE							

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

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

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

4. Reassign the controller by choosing Yes.

A confirmation message is displayed:

NOTICE: Change made to this setting will NOT take effect until the controller is RESET. Prior to resetting the controller, operation may not proceed normally. Do you want to reset the controller now?

5. Choose Yes to reset the controller.

5.5.12 Creating or Changing a Logical Drive Name (Optional)

You can create a name for a logical drive. This logical drive name is used only in RAID firmware administration and monitoring and does not appear anywhere on the host. You can also edit this drive name.

You can create a logical drive name after the logical drive is created:

1. Choose “view and edit Logical drives” from the Main Menu.
2. Select the logical drive and press Return.
3. Choose “logical drive Name.”

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

4. Type the name you want to give the logical drive and press Return to save the name.

5.5.13 Partitioning a Logical Drive (Optional)

You can divide a logical drive into several partitions, or use the entire logical drive as a single partition. You can configure up to 128 partitions for each logical drive.

For guidelines on setting up 1024 LUNs, see [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)”](#) on page 5-47.



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

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

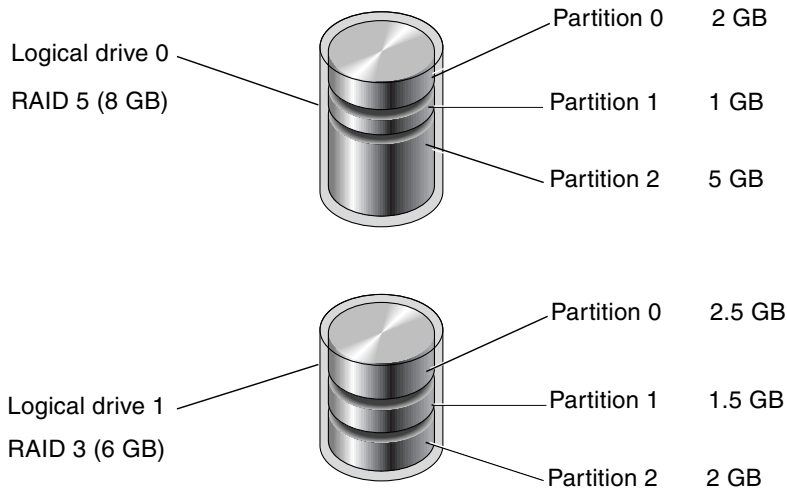


FIGURE 5-5 Partitions in Logical Configurations

To partition a logical drive, perform the following steps.

1. Choose “view and edit Logical drives” from the Main Menu.

```

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

```

2. Select the logical drive you want to partition, and press Return.

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

3. Choose “Partition logical drive.”

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME	
U	P0	27D0811C	NA	RAID0	206856	GOOD	S	6	-	0		
U	S	View scsi drives Delete logical drive Partition logical drive logical drive Name logical drive Assignments Expand logical drive add Scsi drives cOpy and replace drive					GOOD	S	6	-	0	
U	P						GOOD	S	6	-	0	
U	S						GOOD	S	6	-	0	
U	6			NONE								
U	7			NONE								

This message is displayed:

```

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

Continue Partition Logical Drive?

```

Note – Logical volumes are unsuited to some modern configurations, such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information see [“Configuration Overview”](#) on page 5-7.

- Choose **Yes** and press **Return** to confirm that you want to partition the logical drive if you do not want to include it in a logical volume.

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

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

- Select from the list of undefined partitions and press **Return**.
- Type the desired size for the selected partition and press **Return**.

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

Partition Size <MB>: 3000

A warning prompt is displayed:

This operation will result in the loss of all data on the partition.
Partition Logical Drive?



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

7. Choose Yes to confirm.

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

Q	LG	ID	LU	RAID	Size<MB>	Partition	Offset<MB>	Size<MB>	NAME
U	P0	B61E5AB	NA	RAID5	30000	0	0	3000	
U	1			NONE		1	3000	27000	
U	2			NONE		2			
U	3			NONE		3			

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

You can create up to 128 partitions per logical drive, with a total number of partitions not to exceed 1024 partitions per RAID array.

Note – When you modify a partition or logical drive size, you must reconfigure all host LUN mappings. All the host LUN mappings are removed with any change to partition capacity. See [“First Steps in Mapping a Partition to a LUN”](#) on page 5-48.

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

5.6 Mapping Logical Drive Partitions to Host LUNs


The next step is to map each storage partition as one system drive (host ID/LUN). The host adapter recognizes the system drives after reinitializing the host bus.

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

A FC channel can connect up to 16 IDs in loop mode.

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

- The FC ID is like a cabinet, and the drawers are the LUNs (LUN is short for *logical unit number*).
- Each cabinet (ID) can have up to 32 drawers (LUNs).
- Data can be stored in one of the LUNs of the FC ID. Most FC host adapters treat a LUN like another FC device.
- The maximum number of LUNs that can be created for a RAID array is 1024 in loop mode.
- To create a total of 1024 LUNs, see [“Planning for 1024 LUNs \(Optional, Loop Mode Only\)”](#) on page 5-47.



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

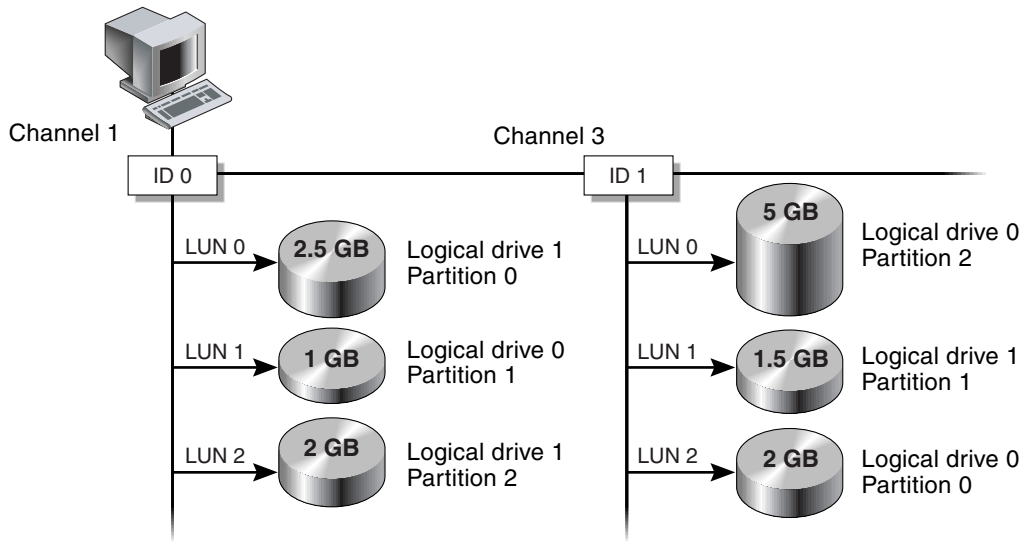


FIGURE 5-6 Mapping Partitions to Host ID/LUNs

5.6.1 Planning for 1024 LUNs (Optional, Loop Mode Only)

If you want to create 1024 LUNs, which is the maximum number of storage partitions that can be mapped for a RAID array, you must map 32 IDs to the array's channels. There are several ways you can meet this requirement. For example, you can set up the following configuration:

- Keep the four default host channels (CH 0, 1, 4, and 5).
- Create eight host IDs per host channel (four primary controller IDs and four secondary controller IDs per host channel) for a total of 32 host IDs. See [“Editing and Creating Additional Host IDs \(Optional\)”](#) on page 5-25.
- Create eight logical drives. See [“Creating Logical Drives \(Optional\)”](#) on page 5-31.

- Partition each logical drive into 128 partitions (8 times 128 = 1024). Map the 1024 partitions to the 32 host IDs. See “Partitioning a Logical Drive (Optional)” on page 5-42 and “Mapping Logical Drive Partitions to Host LUNs” on page 5-45.

TABLE 5-10 Configuration for 1024 LUNs

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

5.6.2 First Steps in Mapping a Partition to a LUN

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

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

```

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

```

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

2. Choose the desired channel and ID to which the logical drive will be mapped, and press Return.

```

----- < Main Menu > -----
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
v
v CHL 0 ID 40 <Primary Controller>
v CHL 1 ID 42 <Secondary Controller>
v CHL 4 ID 44 <Primary Controller>
s CHL 5 ID 46 <Secondary Controller>
v Edit Host-ID/WWN Name List
v

```


5. Select the desired logical drive (LD) and press Return.

The partition table is displayed.

Quick installation view and edit Logical drives view and edit logical Volumes											
LUN	LU/LD	DRU	Partition	Size(MB)	RAID						
0											
LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME	
P0	4DB84961	NA	RAID5	4000	GOOD	S	3	0	0		
CHL	Partition	Offset(MB)	Size(MB)								
0	0	0	4000								
5											
6											
7											

6. Select the desired partition and press Return.

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

7. Select the mapping option that is appropriate for your network and proceed with one of the following procedures.

- “Map Host LUN” is used when multiple hosts are *not* on the same loop. To use this option, see [“Using the Map Host LUN Option” on page 5-50](#).
- “Create Host Filter Entry” is used when multiple hosts share the same loop, can view all the drives, and need to be filtered so that a host sees only the logical drives that are exclusive to its use. To use this option, see [“Setting Up Host Filter Entries” on page 5-52](#).

5.6.3 Using the Map Host LUN Option

Each partition must be mapped to a host LUN. The “Map Host LUN” menu option is used when multiple hosts are not on the same loop.

If multiple hosts will share the same loop on the array, use the host filter command and see [“Setting Up Host Filter Entries” on page 5-52](#).

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

1. After the completion of the steps in “First Steps in Mapping a Partition to a LUN” on page 5-48, choose “Map Host LUN” and press Return.



2. Choose Yes to confirm the mapping scheme.

LUN	LU/LD	DRV	Partition	Size(MB)	RAID
0					
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>< Main Menu ></p> <p>Quick installation</p> <p>view and edit Logical drives</p> <p>view and edit logical Volumes</p> <p>view and edit Host luns</p> <p>u</p> <p>u CHL 0 ID 40 <Primary Control</p> <p>u CHL 1 ID 43 <Secondary Contr</p> <p>u CHL 4 ID 44 <Primary Control</p> <p>s CHL 5 ID 46 <Secondary Contr</p> <p>v Edit Host-ID/WWN Name List</p> </div> <div style="width: 50%;"> <p>Map Logical Drive: 0</p> <p>To Partition : 0</p> <p>Channel : 0</p> <p>ID : 40</p> <p>Lun : 0 ?</p> <p>Yes <input type="checkbox"/> No <input checked="" type="checkbox"/></p> </div> </div>					
5					
6					
7					

The partition is now mapped to a LUN.

LUN	LU/LD	DRV	Partition	Size(MB)	RAID
0	LD	0	0	4000	RAIDS
1					
2					
3					
4					
5					
6					
7					

3. Press the Esc key to return to the Main Menu.
4. Repeat Step 1 through Step 3 for each partition until all partitions are mapped to LUNs.
5. On the Main Menu, choose “system Functions → Reset controller,” and choose Yes to confirm to implement the new configuration settings.
6. To verify unique mapping of each LUN (unique LUN number, unique DRV number, or unique Partition number), choose the “view and edit Host luns” command and press Return.

7. Select the appropriate controller and ID and press Return to review the LUN information.

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

5.6.4 Setting Up Host Filter Entries

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

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




FIGURE 5-7 Example of LUN Filtering

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

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

Note – It is possible to see somewhat different information when a Fabric switch queries the WWN of the Sun StorEdge 3510 FC Array. When the RAID controller does a Fibre Channel Fabric Login to a switch, during the Fabric login process the switch obtains the WWN of the RAID controller. This WWN presented by the RAID controller is a Dot Hill Systems Corporation WWN, so the switch displays this company name. When the switch issues an inquiry command to a mapped LUN on the array, the switch obtains the company name from the inquiry data of the LUN. In this case, the switch displays Sun StorEdge 3510, which is the inquiry data returned by the RAID controller.

As shown in [FIGURE 5-7](#), when you map LUN 01 to host channel 0 and select WWN1, server A has a proprietary path to that logical drive. All servers continue to see and access LUN 02 and LUN 03 unless filters are created on them.

Prior to using the LUN Filter feature, identify which array is connected to which HBA card, and the WWN assigned to each card. This procedure varies according to the HBA you are using. Refer to the appendix for your host for instructions on identifying the WWN for your host.

5.6.4.1 Creating Host Filter Entries

The “Create Host Filter Entry” command is used when multiple hosts share the same loop, can view all the drives, and need to be filtered so that a host sees only the logical drives that are exclusive to its use.

“Map Host LUN” is used when multiple hosts are *not* on the same loop. To use this option, see [“Using the Map Host LUN Option” on page 5-50](#).

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

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

1. After the completion of the steps in the procedure “First Steps in Mapping a Partition to a LUN” on page 5-48, choose “Create Host Filter Entry → Add from current device list.”

Main Menu					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
2					
3					
4					
5					
6					
7					

This step automatically performs a discovery of the attached HBAs. Alternatively, you can add them manually.

2. From the device list, choose the server WWN number for which you are creating a filter and press Return.

Main Menu					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
2					
3					
4					
5					
6					
7					

3. At the confirmation screen, choose Yes and press Return.

Main Menu					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
2					
3					
4					
5					
6					
7					

- Review the filter configuration screen. Make any changes necessary by using the arrow keys to choose an item and press Return.

< Main Menu >						
Quick installation						
view and edit Logical drives						
view and edit logical Volumes						
view and edit Host luns						
v	CHL 0 ID 40 <Primary Contro					
v	CHL 1 ID 42 <Secondary Cont					
v	CHL 4 ID 44 <Primary Contro					
s	CHL 5 ID 46 <Secondary Cont					
v	Edit Host-ID/WWN Name List					

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

Map Host LUN			
Create Host Filter Entry			

Logical Drive 0 Partition 0	
Host-ID/WWN - 0x000000000323542	
Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF	
Filter Type - Include	
Access Mode - Read/Write	
Name - Not Set	

- To edit the WWN, use the arrow keys to choose "Host-ID/WWN" and press Return.

< Main Menu >						
Quick installation						
view and edit Logical drives						
view and edit logical Volumes						
view and edit Host luns						
v	CHL 0 ID 40 <Primary Contro					
v	CHL 1 ID 42 <Secondary Cont					
v	CHL 4 ID 44 <Primary Contro					
s	CHL 5 ID 46 <Secondary Cont					
v	Edit Host-ID/WWN Name List					

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

Map Host LUN			
Create Host Filter Entry			

Logical Drive 0 Partition 0	
Host-ID/WWN - 0x000000000323542	
Host-ID/WWN:	

- Make the desired changes, and press Return.



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

- To edit the WWN Mask, use the arrow keys to choose "Host-ID/WWN Mask" and press Return.

< Main Menu >						
Quick installation						
view and edit Logical drives						
view and edit logical Volumes						
view and edit Host luns						
v	CHL 0 ID 40 <Primary Contro					
v	CHL 1 ID 42 <Secondary Cont					
v	CHL 4 ID 44 <Primary Contro					
s	CHL 5 ID 46 <Secondary Cont					
v	Edit Host-ID/WWN Name List					

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

Map Host LUN			
Create Host Filter Entry			

Logical Drive 0 Partition 0	
Host-ID/WWN - 0x000000000323542	
Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF	
Host-ID/WWN Mask: _	

8. To change the filter setting, use the arrow keys to choose "Filter Type -" and press Return.
9. At the confirmation screen, choose Yes to exclude or No to include the Host-ID/WWN selection, and press Return.

< Main Menu >					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns v v CHL 0 ID 40 <Primary Contro v CHL 1 42 <Secondary Cont v CHL 4 44 <Primary Contro s CHL 5 46 <Secondary Cont v Edit Host-ID/WWN Name List					
Map Host LUN Create Host Filter Entry					
4	Logical Drive 0 Partition 0 Host-ID/WWN - 0x0000000000323542 Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF				
5	Filter Type - Include				
6	A N Set Filter Type to Exclude ?				
7	Yes No				

10. To change the access mode that assigns Read-Only or Read/Write privileges, use the arrow keys to choose "Access mode -" and press Return.
11. At the confirmation screen, choose Yes and press Return.

< Main Menu >					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns v v CHL 0 ID 40 <Primary Contro v CHL 1 ID 42 <Secondary Contro v CHL 4 ID 44 <Primary Contro s CHL 5 ID 46 <Secondary Contro v Edit Host-ID/WWN Name List					
Map Host LUN Create Host Filter Entry					
4	L Set Access Mode to Read-Only ? H H Yes No F				
5	Access Mode - Read/Write				
6	Name - Not Set				
7					

12. To set a name, use the arrow keys to choose "Name -" and press Return.

< Main Menu >					
LUN	LU/LD	DRU	Partition	Size(MB)	RAID
0	LD	0	0	150000	RAID5
1					
Quick installation view and edit Logical drives view and edit logical Volumes view and edit Host luns v v CHL 0 ID 40 <Primary Contro v CHL 1 ID 42 <Secondary Contro v CHL 4 ID 44 <Primary Contro s CHL 5 ID 46 <Secondary Contro v Edit Host-ID/WWN Name List					
Map Host LUN Create Host Filter Entry					
4	L H Name :mars 542 H FFFFFFFF F A				
5					
6	Name - Not Set				
7					

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

- Verify all settings and press Esc to continue.

< Main Menu >						
Quick installation						
view and edit Logical drives						
view and edit logical Volumes						
view and edit Host Luns						
v	CHL 0 ID 40 <Primary Contro					
v	CHL 1 ID 42 <Secondary Cont					
v	CHL 4 ID 44 <Primary Contro					
s	CHL 5 ID 46 <Secondary Cont					
v	Edit Host-ID/WWN Name List					
	LUN	LU/LD	DRU	Partition	Size(MB)	RAID
	0	LD	0	0	150000	RAID5
	1					
	Map Host LUN					
	Create Host Filter Entry					
	4	Logical Drive 0 Partition 0				
		Host-ID/WWN - 0x0000000000323542				
		Host-ID/WWN Mask- 0xFFFFFFFFFFFFFFF				
		Filter Type - Exclude				
		Access Mode - Read/Write				
		Name - mars				
	5					
	6					
	7					

Note – Unlike most firmware operations where you must complete each entry individually and repeat the procedure if you want to perform a similar operation, you can add multiple WWNs to your list before you actually create the host filter entry in [Step 15](#). Be sure to follow the instructions carefully.

- At the confirmation screen, choose Yes and press Return.

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

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

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

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

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

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

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

5.6.5 Creating Device Files for the Solaris Operating Environment

- To create device files for the newly mapped LUNs on the host in the Solaris 8 operating environment and Solaris 9 operating environment, type:

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

- To display the new LUNs, type:

```
# format
```


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

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


5.6.6 Saving Configuration (NVRAM) to a Disk

You can back up your controller-dependent configuration information. Use this function to save configuration information whenever a configuration change is made.


The logical configuration information is stored within the logical drive.

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

1. Choose “system Functions → Controller maintenance → Save nvram to disks.”



2. Choose Yes to confirm.



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

To restore the configuration, see [“Restoring Your Configuration \(NVRAM\) From a File”](#) on page 7-24.

5.7 Using Software to Monitor and Manage the Sun StorEdge 3510 FC Array

This section describes the software management tools available for monitoring and managing the Sun StorEdge 3510 FC array with in-band connections.

The following software management tools are provided on the Sun StorEdge 3000 Family Professional Storage Manager CD, provided with your array. The Sun StorEdge 3000 Family Documentation CD provides the related user guides.

- **Sun StorEdge Configuration Service.** Provides centralized storage configuration, maintenance, and monitoring functions. Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for in-band setup procedures.
- **Sun StorEdge Diagnostic Reporter.** Provides event monitoring and notification. Refer to the *Sun StorEdge Diagnostic Reporter User's Guide* for information.
- **Sun StorEdge Command-Line Interface (CLI).** A command-line interface utility that provides script-based management. Refer to the *Sun StorEdge CLI User's Guide* for CLI information.

For details on how to install Sun StorEdge Configuration Service, Sun StorEdge Diagnostic Reporter, or the Sun StorEdge CLI software, refer to the *Sun StorEdge 3000 Family Software Installation Manual*.

The Sun StorEdge 3000 Family Documentation CD provides the related user guides, with detailed installation and configuration procedures for Sun StorEdge Configuration Service and Sun StorEdge Diagnostic Reporter.

5.7.1 Other Supported Software

Multipathing for the Sun StorEdge 3510 FC Array is provided by Sun StorEdge Traffic Manager software. Multipathing software is required when you have multiple connections from a server to an array (directly or through a switch), want to avoid a single point of failure, and are setting up a configuration with redundant pathing. Multipathing software makes multiple paths between the server and the storage system and provides full services on each path for path failover.

Refer to the appendix for your host and to the *Sun StorEdge 3510 FC Array Release Notes* for information about which versions of Sun StorEdge Traffic Manager software are supported on your platform.

For information about additional supported or provided software, also refer to the *Sun StorEdge 3510 FC Array Release Notes*.

5.7.2 Enabling VERITAS DMP

To enable VERITAS Dynamic Multi-Pathing (DMP) support on VERITAS Volume Manager in the Solaris operating environment, perform the following steps.

Note – To see instructions for enabling VERITAS DMP on other supported platforms, refer to your VERITAS user documentation.

1. Configure at least two channels as host channels (channels 1 and 3 by default) and add additional host IDs, if needed.
2. Connect host cables to the I/O host ports in Step 1.
3. Map each LUN to two host channels to provide dual-pathed LUNs.
4. Add the correct string to `vxddladm` so VxVM can manage the LUNs as a multipathed JBOD.

```
# vxddladm addjbod vid=SUN pid="StorEdge 3510"
# vxddladm listjbod
VID          PID          Opcode   Page     Code    Page Offset  SNO length
=====
SEAGATE ALL          PIDs      18      -1       36          12
SUN          StorEdge 3510     18      -1       36          12
```

5. Reboot the hosts. System reboot is required to implement these changes.

Note – JBOD arrays require a license from VERITAS in order to enable any of its advanced features. Refer to VERITAS Volume Manager Release Notes or contact VERITAS Software Corporation for licensing terms and information.

5.7.3 The VERITAS Volume Manager ASL

VERITAS has provided an Array Software Library (ASL) that must be installed on the same host system as the Volume Manager 3.2 or 3.5 software to enable the software to recognize the Sun StorEdge 3510 FC array. For the procedure to download the ASL and the accompanying installation guide for the Sun StorEdge 3510 FC array from Sun's Download Center, refer to the *Sun StorEdge 3510 FC Array Release Notes*.

Checking LEDs

This chapter describes the front and back panel LEDs, which give the clear status of the operation of all drives and modules. Topics covered in this chapter are:

- “LEDs When Array Is First Powered On” on page 6-1
- “Front Panel LEDs” on page 6-1
- “Back Panel LEDs” on page 6-5

For troubleshooting flowcharts related to LEDs, refer to “[Troubleshooting Flowcharts](#)” on page 8-15.

6.1 LEDs When Array Is First Powered On

With the array powered up but not connected to a server, you should see the LED conditions described in [TABLE 6-1](#).

TABLE 6-1 Front Panel LED Status When Array Is First Powered On

Drive LEDs	Solid green
Chassis ear LEDs	Solid green

6.2 Front Panel LEDs

The drive LEDs are located between the rows of drives on the front panel, as shown in [FIGURE 6-1](#). The system operation LEDs are located on the right ear of the chassis.




FIGURE 6-1 Front Panel LEDs

FIGURE 6-2 shows the front panel LEDs and the Reset button. Use a paperclip to push the Reset button to silence a failed component alarm. See [Section 7.2, “Silencing Audible Alarms”](#) on page 7-5 for more information about silencing audible alarms.




FIGURE 6-2 Chassis Ear LEDs and Reset Button on Front Panel

The following table lists the front panel LEDs.

TABLE 6-2 Front Panel LEDs

LED	LED Color	Description
Drive	Solid green	Good: Drive power-up and spin-up OK.
	Blinking green	Good: Drive activity.
	Solid amber	Failed: Drive failure.
Power (Light bulb icon) Monitors the DC output voltage within tolerance specification. Overcurrent protection shutting down any voltage output is also displayed. Voltage thresholds: +5 VDC +/- .25 VDC +12 VDC +/- .6 VDC Current thresholds: +5 VDC 35A +12 VDC 25A	Solid green	Good: Power supply good.
	Solid amber	Failed: One or more output voltages out-of-range.
Fan (Fan icon) Monitors the fan speed within nominal operational RPM specification of 5000 RPM.	Solid green	Good: Both fans rotating at 4000 RPM or higher.
	Solid amber	Faulty/Failed: One or both fans rotating a less than 4000 RPM.
Temp (Thermometer icon) Monitors the temperature level and indicates violations of internal temperature thresholds of 55C and 60C.	Solid green	Good: Under temperature threshold of 55C.
	Solid amber	Failed: Temperature threshold equal to or over 55C.
	Blinking amber	Failed: Temperature threshold equal to or over 60C. Blinking frequency of 4 Hz +/- 1 Hz.
Event (Caution icon) Indicates any abnormal or failure event in the I/O board.	Solid green	Normal operation of I/O board.
	Solid amber	Failed I/O board.
	Blinking amber	Indicates that the version of the SES firmware or associated hardware PLD code on one controller does not match that on the other controller.

Note – To test that the LEDs work, using a paperclip, press and hold the Reset button for 5 seconds. All the LEDs should change from green to amber when you perform this test. Any LED that fails to light indicates a problem with the LED. When you release the Reset button, the LEDs return to their initial state.

6.2.1 Correcting SES or PLD Firmware Version Conflicts

From time to time, firmware upgrades are made available as patches that you can download from SunSolve™ Online, located at:

<http://sunsolve.sun.com>

Each patch applies to a particular piece of firmware, including the firmware programmed into SES and PLD chips included in your controller.

SunSolve has extensive search capabilities that can help you find these patches, as well as regular patch reports and alerts to let you know when firmware upgrades and other patches become available. In addition, SunSolve provides reports about bugs that have been fixed in patch updates.

Each patch includes an associated README text file that provides detailed instructions about how to download and install that patch. But, generally speaking, all firmware downloads follow the same steps:

- Locating the patch on SunSolve that contains the firmware upgrade you want
- Downloading the patch to a location on your network
- Using your array software (SSCS or `sscli`[1M]) or array firmware, in some cases, to “flash” the firmware to the device it updates

Refer to the release notes for your array for patch numbers of firmware and other patches available for your array at the time of release.

When you replace an I/O controller, the new controller might have a version of SES or PLD firmware different from the other controller in your array. If this mismatch occurs, when you install a controller you hear an audible alarm and see a blinking amber Event light.

To synchronize the SES firmware and hardware PLD versions, you need to download new SES firmware through Sun StorEdge Configuration Service software or the command-line interface (CLI).

If you have not installed this software, you need to install it from the software CD that shipped with your array. Refer to the *Sun StorEdge 3000 Family Configuration Service User's Guide* for your array to see instructions for downloading firmware for

devices, or refer to the `sccli(1M)` man page for similar instructions using the CLI. Refer to the release notes for your array for instructions about where to obtain the firmware that you need to download.

When you open Sun StorEdge Configuration Service software or the CLI and connect to the array, an error message alerts you to the mismatched version problem.

6.3 Back Panel LEDs

The back panel LED colors indicate the conditions described in the following figures and tables.

Note – Although an amber light on the back panel often indicates a faulty component, a steady amber light on the Ethernet link LED indicates normal Ethernet activity. See [TABLE 6-3](#) for more information.

6.3.1 I/O Controller Module LEDs

[FIGURE 6-3](#) shows the I/O controller module and its LEDs on the back panel.




FIGURE 6-3 I/O Controller Module and Battery Module LEDs

FIGURE 6-4 shows the I/O expansion module and its LEDs.




FIGURE 6-4 I/O Expansion Module for an Expansion Unit

The I/O controller module LEDs and their color definitions are shown in TABLE 6-3.

TABLE 6-3 I/O Controller Module and Battery Module LEDs

LED	Purpose	LED Color Definition
Battery	Status of battery	Solid green – Battery charged Blinking green – Battery charging Solid amber – Battery failed
Activity	I/O activity of host and disk ports	Off – Not busy (no I/O activity) Blinking green – Busy (active I/O)
Cache	Status of memory cache	Off – Clean cache Blinking green – Dirty cache; indicates if data in memory and is not written to disk yet
Ethernet link (active controller)	Status of Ethernet link	Solid amber – Active link Off – Inactive connection
Ethernet link (inactive controller)	Status of Ethernet link	Off – Active connection on inactive controller or on active connection
Ethernet active	Status of Ethernet activity	Flashing green – Busy

TABLE 6-3 I/O Controller Module and Battery Module LEDs (*Continued*)

LED	Purpose	LED Color Definition
RAID controller	Status of controller on I/O controller module	Blinking green – Good (primary controller) Solid green – Good (secondary controller) Solid amber – Failed RAID controller or I/O module
SFP link (L)	SFP link status	Solid green – Active good FC Connection Off – Empty or failed FC connection
SFP speed (S)	SFP speed status	Solid green – 2 Gbit Off – 1 Gbit

Note – The Sun StorEdge 3510 FC array’s Ethernet link LED differs from that found on a Sun StorEdge 3310 SCSI array. A connected port’s Ethernet link LED glows solid amber only on the active controller, because a FC port is considered active only if its controller is active. On an inactive controller, a port’s Ethernet link LED is off whether it is connected or not. The absence of a solid amber light on a connected port’s Ethernet link LED for a connected port indicates that this controller is not the active controller.

Note – An SFP link status (L) LED is off if there is no connection, or a failed connection, to this SFP.

Note – To test that the LEDs work, using a paperclip, press and hold the Reset button for 5 seconds. All the LEDs should change from green to amber when you perform this test. Any LED that fails to light indicates a problem with the LED. When you release the Reset button, the LEDs return to their initial state.

6.3.2 Power Supply and Fan Module LEDs

TABLE 6-4 Power Supply LEDs

Purpose	LED	LED Color Definition
Monitors the DC output voltage within tolerance specification. Overcurrent protection shutting down any voltage output is also displayed. Voltage thresholds: +5 VDC +/- .25 VDC +12 VDC +/- .6 VDC Current thresholds: +5 VDC 35A +12 VDC 25A	Solid green Solid amber	Power supply and fans are good. Failed: One or more output voltages out-of-range or fan speed is under 3150 RPM.

The following figure shows the AC power supply and fan module.




FIGURE 6-5 AC Power Supply and Fan Module

The following figure shows the DC power supply and fan module.




FIGURE 6-6 DC Power Supply and Fan Module

Maintaining Your Array

This chapter covers the following maintenance and troubleshooting topics:

- “Introducing Key Screens and Commands” on page 7-2
 - “Controller Firmware Initial Screen” on page 7-2
 - “Main Menu” on page 7-4
 - “Quick Installation (Reserved)” on page 7-4
- “Silencing Audible Alarms” on page 7-5
- “Deleting Logical Drives” on page 7-6
- “Checking Status Windows” on page 7-8
 - “Logical Drive Status Table” on page 7-9
 - “Physical Drive Status Table” on page 7-11
 - “Channel Status Table” on page 7-13
 - “Controller Voltage Temperature Status” on page 7-15
 - “Viewing SES Status” on page 7-18
 - “Viewing Event Logs on the Screen” on page 7-22
- “Restoring Your Configuration (NVRAM) From a File” on page 7-24
- “Upgrading Firmware” on page 7-26
 - “Patch Downloads” on page 7-27
 - “Controller Firmware Upgrade Features” on page 7-27
 - “Installing Firmware Upgrades” on page 7-28
 - “Installing Controller Firmware Upgrades From the Firmware Application” on page 7-29
- “Replacing the Front Bezel and Ear Caps” on page 7-30

7.1 Introducing Key Screens and Commands

This section introduces the initial RAID controller firmware screen and Main Menu.

7.1.1 Controller Firmware Initial Screen

You see the following initial controller screen when you first access the RAID controller firmware via the controller COM port or Ethernet port.

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




TABLE 7-1 Components of the Controller Firmware Screen

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

TABLE 7-1 Components of the Controller Firmware Screen (Continued)

Component	Description
(VT100 mode)	Enters the Main Menu and operates in VT100 mode.
PC graphic (ANSI+color mode)	Enters the Main Menu and operates in ANSI color mode.
Show transfer rate+show cache status	Press Return on this item to show the cache status and transfer rate.

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

LG	ID	LU	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
P0	488FFBB8	NA	RAID5	103428	GOOD				7	B	4	0	0	
P1	34E	Drive Copying												
2		-----												
3		87% Completed												

Events showing full descriptive titles for the progress indicator include:

- Drive Copying
- Flash Erasing
- Flash Programming

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

TABLE 7-2 Progress Indicator Prefix Meanings

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

Note – For information about the battery status indicator, see ["Battery Status" on page 5-4.](#)

7.1.2 Main Menu

After you have selected the VT100 terminal emulation display mode and pressed Return on the initial screen, the Main Menu is displayed.

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

Use the following keys to navigate within the application:

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

7.1.3 Quick Installation (Reserved)

This menu option is not used in normal operation. It is reserved for special use in special situations, and only when directed by technical support.



Caution – Do not use this menu option unless directed by technical support. Using it will result in the loss of your existing configuration and all data you have on the devices.

7.2 Silencing Audible Alarms

An audible alarm indicates that either a component in the array has failed or a specific controller event has occurred. Error conditions and controller events are reported by event messages and event logs. Component failures are also indicated by LED activity on the array.

Note – It is important to know the cause of the error condition because how you silence the alarm depends on the cause of the alarm.

To silence the alarm, perform the following steps:

1. **Check the error messages, event logs, and LED activity to determine the cause of the alarm.**

Component event messages include but are not limited to the following terms:

- SES/PLD firmware mismatch
- Temperature
- Cooling element
- Power supply
- Battery
- Fan
- Voltage sensor



Caution – Be particularly careful to observe and rectify a temperature failure alarm. If you detect this alarm, shut down the controller and the server as well if it is actively performing I/O operations to the affected array. Otherwise, system damage and data loss can occur.

See [“Failed Component Alarm Codes”](#) on page C-1 for more information about component alarms.

Controller event messages include but are not limited to the following terms:

- Controller
- Memory
- Parity
- Drive SCSI Channel
- Logical drive
- Loop connection

Refer to the “Event Messages” appendix in the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for more information about controller events.

2. Depending on whether the cause of the alarm is a failed component or a controller event and which application you are using, silence the alarm as specified in the following table.

TABLE 7-3 Silencing the Alarm

Cause of Alarm	To Silence Alarm		
	Firmware Application	Sun StorEdge Configuration Service	Command-Line Interface
Failed Component Alarms	Use a paperclip to push the Reset button on the right ear of the array.		
Controller Event Alarms	From the Main Menu, choose "system Functions → Mute beeper." Refer to the <i>Sun StorEdge 3000 Family RAID Firmware User's Guide</i> for more information.	Refer to "Updating the Configuration" in the <i>Sun StorEdge Configuration Service User's Guide</i> for information about the "Mute beeper" command.	Run <code>mute [controller]</code> . Refer to the <i>Sun StorEdge 3000 Family CLI User's Guide</i> for more information.

Note – Pushing the Reset button has no effect on controller event alarms and muting the beeper has no effect on failed component alarms.

7.3 Deleting Logical Drives

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



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

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

To first unmap and then delete a logical drive, perform the following steps:

1. Choose the “view and edit Host luns” from the Main Menu.
Existing logical drive mappings are displayed.
2. Select the existing logical drive you want to unmap and press Return.
A menu of host LUNs is displayed.
3. Select the host LUN that you want to unmap and press Return.
A confirmation message asks if you want to unmap the host LUN you have selected.
4. Choose Yes to unmap the logical drive.
5. Repeat [Step 3](#) and [Step 4](#) to unmap all LUNS that are mapped to the logical drive you want to delete.
6. From the Main Menu, choose “view and edit Logical drives.”

```

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

```

7. Select the logical drive you unmapped and want to delete, and press Return.

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

8. Choose "Delete logical drive."

Q	LG	ID	LU	RAID	Size(MB)	Status	O	#LN	#SB	#FL	NAME
U	P0	22D0811C	NA	RAID0	206856	GOOD	S	6	-	0	
U	S	View scsi drives				GOOD	S	6	-	0	
U	S	Delete logical drive									
U	P	Partition logical drive				GOOD	S	6	-	0	
U	S	logical drive Name									
U	S	logical drive Assignments				GOOD	S	6	-	0	
U	S	Expand logical drive									
U	S	add Scsi drives									
U	S	copy and replace drive									
	6			NONE							
	7			NONE							

A warning notice is displayed asking if you are certain you want to delete the logical drive and its data.

9. Choose Yes to confirm.

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

7.4 Checking Status Windows

The status windows used to monitor and manage the array are described in the following sections:

- "Logical Drive Status Table" on page 7-9
- "Physical Drive Status Table" on page 7-11
- "Channel Status Table" on page 7-13
- "Controller Voltage Temperature Status" on page 7-15
- "Viewing SES Status" on page 7-18
- "Viewing Event Logs on the Screen" on page 7-22

7.4.1 Logical Drive Status Table

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

```

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

The status of all logical drives is displayed.

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

TABLE 7-4 shows definitions and values for logical drive parameters.

TABLE 7-4 Parameters Displayed in the Logical Drive Status Window

Parameter	Description
LG	Logical drive number P0: Logical drive 0 of the primary controller where P = primary controller and 0 = logical drive number S1: Logical drive 1 of the secondary controller where S = secondary controller and 1 = logical drive number
ID	Logical drive ID number (controller-generated)
LV	The logical volume to which this logical drive belongs. NA indicates no logical volume.

TABLE 7-4 Parameters Displayed in the Logical Drive Status Window *(Continued)*

Parameter	Description
RAID	RAID level
SIZE (MB)	Capacity of the logical drive in megabytes
Status	Logical drive status
INITING	The logical drive is now initializing.
INVALID	The logical drive was improperly created or modified. For example, the logical drive was created with "Optimization for Sequential I/O," but the current setting is "Optimization for Random I/O."
GOOD	The logical drive is in good condition.
DRV FAILED	A drive member failed in the logical drive.
FATAL FAIL	More than one drive member in a logical drive has failed.
REBUILDING	The logical drive is rebuilding.
DRV ABSENT	One of the disk drives cannot be detected.
INCOMPLETE	Two or more member disk drives in this logical drive have failed.
O	Indicates the performance optimization set when the logical drive was initialized. This cannot be changed after the logical drive is created. S Optimization for Sequential I/O R Optimization for Random I/O
#LN	Total number of drive members in this logical drive
#SB	Number of standby drives available for the logical drive. This includes local spare and global spare disk drives available for the logical drive.
#FL	Number of failed disk drive members in the logical drive
Name	Logical drive name (user configurable)

To handle failed, incomplete, or fatal failure status, see ["Identifying a Failed Drive for Replacement"](#) on page 8-7 and ["Recovering From Fatal Drive Failure"](#) on page 8-11.

7.4.2 Physical Drive Status Table

To check and configure physical drives, from the Main Menu choose “view and edit scsi Drives” and press Return.

```

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

The screen displays the status of all physical drives.

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

TABLE 7-5 Parameters Displayed in the Drive Status Window

Parameters	Description
Slot	Slot number of the drive
Ch1	Channel of the connected drive
ID	ID of the drive
Size (MB)	Drive capacity in megabytes
Speed	xxMB Maximum synchronous transfer rate of this drive. Async The drive is using asynchronous mode.
LG_DRV	x The drive is a physical drive member of logical drive x.
Status	GLOBAL The drive is a global spare drive.

TABLE 7-5 Parameters Displayed in the Drive Status Window (Continued)

Parameters	Description
INITING	The drive is initializing.
ON-LINE	The drive is in good condition.
REBUILD	The drive is rebuilding.
STAND-BY	Local spare drive or global spare drive. The local spare drive's LG_DRV column shows the logical drive number. The global spare drive LG_DRV column shows "Global."
NEW DRV	The new drive has not been configured to any logical drive or as a spare drive.
USED DRV	The drive was previously configured as part of a logical drive from which it has been removed; it still contains data from this logical drive.
FRMT DRV	The drive has been formatted with reserved space allocated for controller-specific information.
BAD	Failed drive.
ABSENT	Drive slot is not occupied or the drive is defective and cannot be detected.
MISSING	Drive once existed, but is now missing.
SB-MISS	Spare drive missing.
Vendor and product ID	Vendor and product model information of the drive.

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

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

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

To handle BAD drives, refer to ["Identifying a Failed Drive for Replacement" on page 8-7](#).

If two drives show BAD and MISSING status, see ["Recovering From Fatal Drive Failure" on page 8-11](#).

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

Note – When power is on, the controller scans all hard drives that are connected through the drive channels. If a hard drive was connected after the controller completes initialization, use the “Scan scsi drive” submenu option after you have selected a drive to let the controller recognize the newly added hard drive and configure it as a member of a logical drive.

7.4.3 Channel Status Table

To check and configure channels, from the Main Menu choose “view and edit Scsi channels,” and press Return.

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

The screen displays the status of all channels for this controller.

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

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
1	Host	40	NA	2 GHz	Serial	F	NA	2 GHz	Serial
1	Host	NA	42	2 GHz	Serial	F	NA	2 GHz	Serial
2<3;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
3<2;D>	DRU+RCC	14	15	2 GHz	Serial	F	NA	2 GHz	Serial
4	Host	44	NA	2 GHz	Serial	F	NA	2 GHz	Serial
5	Host	NA	46	2 GHz	Serial	F	NA	2 GHz	Serial

Note – Each controller has a separate RS232 port, as well as an Ethernet chip. This architecture ensures continuous communication should a controller fail. Since the connection is established to only one controller (even when the array is in redundant mode), the CurSyncClk and CurWid settings are displayed for that individual controller. Therefore, if a user maps one LUN to the primary controller, and another LUN to a secondary controller, only the established connection to that individual controller is displayed through the serial and Ethernet port menu.



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

TABLE 7-6 Parameters Displayed in the Channel Window

Parameters	Description
Chl	Channel's ID.
Mode	Channel mode.
	RCC Redundant controller communication channel.
	Host The channel is functioning as a host channel.
	DRV The channel is functioning as a drive channel.
PID	Primary controller's ID mapping:
	* Multiple IDs were applied (host channel mode only).
	# The ID for host LUNs mapped to this channel in Host Channel mode. ID for the primary controller in drive channel mode.
	NA No ID applied.
SID	Secondary controller's ID mapping:
	* Multiple IDs (Host Channel mode only).
	# The ID for host LUNs mapped to this channel in host channel mode. ID for the secondary controller in drive channel mode.
	NA No ID applied.
DefSynClk	Default bus synchronous clock:
	<i>n</i> GHz Maximum synchronous transfer rate.
	Async Channel is set for asynchronous transfers.
DefWid	Default bus width: Not applicable for Fibre Channel.
	Serial

TABLE 7-6 Parameters Displayed in the Channel Window (*Continued*)

Parameters	Description	
S	Signal:	
	S	Single-ended
	L	LVD
Term	F	Fibre
	Terminator status:	
	On	Termination is enabled.
CurSynClk	Off	Termination is disabled.
	NA	For a redundant controller communications channel (RCCOM).
	Current bus synchronous clock:	
CurWid	xx.x MHz	The current speed at which the channel is communicating.
	Async.	The channel is communicating asynchronously or no device is detected.
	(empty)	The default bus synchronous clock has changed. Reset the controller for changes to take effect.
CurWid	Current bus width: Not applicable for Fibre Channel.	
	Serial	
	(empty)	The default bus width has changed. Reset the controller for the changes to take effect.

7.4.4 Controller Voltage Temperature Status

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

1. Choose “view and edit Peripheral devices → Controller Peripheral Device Configuration → View Peripheral Device Status.”

The voltage and temperature status of the RAID unit is displayed.

```

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

```

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

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view an
view an
view an
view an
view an
s
v
v
View
Set
Defi
Adju
Cont
F

```

ITEM	VALUE	STATUS
+3.3V	3.384V	Operation Normally
+5V	5.072V	Operation Normally
+12V	12.260V	Operation Normally
CPU Temperature	49.5 (C)	Temperature within Safe Range
Board1 Temperature	52.5 (C)	Temperature within Safe Range
Board2 Temperature	68.0 (C)	Temperature within Safe Range

```

View Peripheral Device Status
Voltage and Temperature Parameters

```

2. Choose "Voltage and Temperature Parameters" and press Return to view or edit the trigger thresholds that determine voltage and temperature status.

```

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

```

3. Select a threshold you want to view or edit and press Return.

```

      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
S
v
v
View
Set Trigger Thresholds for +3.30 Events
Defi Trigger Thresholds for +50 Events
Adju Trigger Thresholds for +120 Events
Cont Trigger Thresholds for CPU Temperature Events
F    Trigger Thresholds for Board Temperature Events
U
Voltage and Temperature Parameters

```

- Repeat this step as many times as necessary to “drill down” to the threshold ranges and triggering events.

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
s
u
v
View
Set Trigger Thresholds for +3.30 Events
Defi T
Adju T Upper Threshold for +3.30 Event - Default<3.60>
Cont T Lower Threshold for +3.30 Event - Default<2.90>
F T
U
Voltage and Temperature Parameters

```

```

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit
view and edit
view and edit
view and edit
view and edit
view and edit
Voltage Range from 3.40 to 3.90
Default Trigger Event: "default"
Disable Trigger Event: "disable"
Input Voltage Trigger Threshold : default
s
u
v
View
Set Trig
Defi T
Adju T Upper Threshold for +3.30 Event - Default<3.60>
Cont T Lower Threshold for +3.30 Event - Default<2.90>
F T
U
Voltage and Temperature Parameters

```

- To edit a trigger or other editable value, backspace over the existing information and replace it.

7.4.5 Viewing SES Status

The array’s SCSI Enclosure Services (SES) processor, located on the I/O module, monitors environmental conditions and is supported by Sun StorEdge Configuration Service and the command-line interface.

For Sun StorEdge 3510 FC JBOD arrays only, both Sun StorEdge Configuration Service and the CLI access the SES processor using device files in `/dev/es`, such as `/dev/es/ses0`, as shown in the following example:

```
# sccli

Available devices:

1. /dev/rdisk/c4t0d0s2 [SUN StorEdge 3310 SN#000280] (Primary)

2. /dev/es/ses0 [SUN StorEdge 3510F D SN#00227B] (Enclosure)
```

To check the status of SES components (temperature sensors, cooling fans, the beeper speaker, power supplies, and slot status), perform the following steps:

1. Choose “view and edit Peripheral devices → View Peripheral Device Status → SES device.”

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
s
v
v
View Peripheral Device Status
S
D
A
C
F
ITEM          STATUS          LOCATION
Redundant Controller    Degraded          Primary
SES Device             Enclosure Device  Channel 2 ID 12
SES Device             Enclosure Device  Channel 2 ID 76
```

A list is displayed of environmental sensors and other hardware components of that SES device.

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view an
view an Enclosure Descriptor Device
view an Cooling element parameters
view an Temperature Sensors ses
view an Voltage sensor
s
v
v
View Power Supply is
S
D
A
C
R
Audible alarm
Nonvolatile cache
SCSI port/transceiver
PBC
ITEM          STATUS          LOCATION
Redundant Controller    Degraded          Primary
SES Device             Enclosure Device  Channel 2 ID 12
SES Device             Enclosure Device  Channel 2 ID 76
```

2. Select an item from the list and press Return to display information about it or see a submenu of its component attributes.

Choosing Overall Status displays the status of the SES device and its operating temperature:

```

< Main Menu >
Quick installation
view and edit Logical drives
view and e
view an Element Descriptor
view an En Overall Status
view an De Element 0
view an Co Element 1
view an Te Element 2
view an Uo Element 3
s v Uview Po Element 4
v S Au Element 5
D No Element 6
A SC Element 7
C R PB Element 8
F Element 9
SES Element 10
SES Element 11

```

STATUS	LOCATION
Degraded	Primary
Enclosure Device	Channel 2 ID 12
Enclosure Device	Channel 2 ID 76

```

< Main Menu >
Quick installation
Status :OK
Temperature :33 <C>
view an De Element 0
view an Co Element 1
view an Te Element 2
s v Uview Po Element 3
v S Au Element 4
D No Element 5
A SC Element 6
C R PB Element 7
F Element 8
SES Element 9
SES Element 10
SES Element 11

```

STATUS	LOCATION
Degraded	Primary
Enclosure Device	Channel 2 ID 12
Enclosure Device	Channel 2 ID 76

3. Select other attributes in which you are interested and press Return to learn more about the SES device.

7.4.6 SES Temperature Sensor Locations (FC Only)

Monitoring temperature at different points within the array is one of the most important SES functions. High temperatures can cause significant damage if they go unnoticed. There are a number of different sensors at key points within the enclosure. The following table shows the location of each of those sensors. The

Element ID corresponds to the identifier shown when you choose “view and edit Peripheral devices → View Peripheral Device Status → SES Device → Temperature Sensors.”

TABLE 7-7 Sun StorEdge 3510 Temperature Sensor Locations

Element ID	Description
0	Drive Midplane Left Temperature Sensor #1
1	Drive Midplane Left Temperature Sensor #2
2	Drive Midplane Center Temperature Sensor #3
3	Drive Midplane Center Temperature Sensor #4
4	Drive Midplane Right Temperature Sensor #5
5	Drive Midplane Right Temperature Sensor #6
6	Upper IOM Left Temperature Sensor #7
7	Upper IOM Left Temperature Sensor #8
8	Lower IOM Temperature Sensor #9
9	Lower IOM Temperature Sensor #10
10	Left PSU Temperature Sensor #11
11	Right PSU Temperature Sensor #12

7.4.7 Identifying Fans

You can view the status of SES components, including the pair of fans located in each fan and power supply module. A fan is identified by the SES Device menus as a cooling element.

Follow these steps to view the status of each fan:

1. Choose “**view and edit Peripheral devices → View Peripheral Device Status → SES Device → Cooling element.**”
2. Select one of the elements (element 0, 1, 2, or 3).

Standard fan speeds are indicated by numbers 1 through 7, indicating speeds in the normal range of 4000 to 6000 RPM. The number 0 indicates that the fan has stopped.

If a fan fails and the Status field does not display the OK value, you must replace the fan and power supply module.

Cooling elements in the status table can be identified for replacement as shown in [TABLE 7-8](#).

TABLE 7-8 Relationship Between Cooling Elements, Fans, and Power Supply Modules

Cooling Element #	Fan # and Power Supply Module #
Cooling Element 0	FAN 0, PS 0
Cooling Element 1	FAN 1, PS 0
Cooling Element 2	FAN 2, PS 1
Cooling Element 3	FAN 3, PS 1

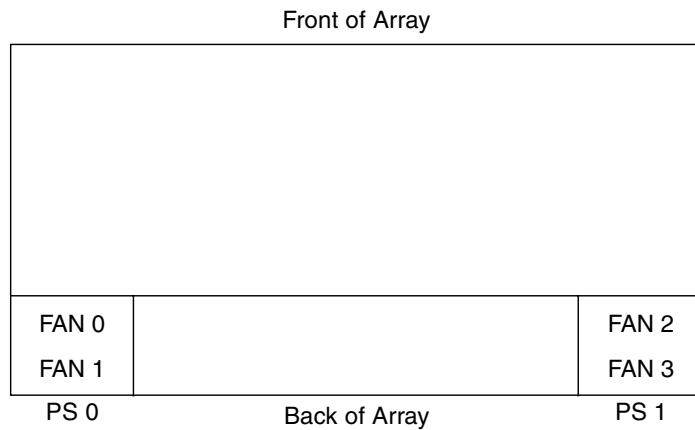


FIGURE 7-1 Cooling Fan Locations

7.4.8 Viewing Event Logs on the Screen

A controller event log records an event or alarm that occurs after the system is powered on. The controller can store up to 1000 event logs. An event log records a configuration or operation event as well as an error message or alarm event.

Note – The SES logic in each array sends messages to the event log, which reports problems and the status of the fans, temperature, and voltage.



Caution – Powering off or resetting the controller automatically deletes all recorded event logs.

1. To view the event logs on screen, choose “view and edit Event logs” on the Main Menu.

```

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

```

A log of recent events is displayed.

Event Logs	
[1106] CHL:2 ID:74 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Wed Dec 25 16:25:49 2002	P
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 17:47:26 2002	P
[1106] CHL:2 ID:4 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 21:33:00 2002	P
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Thu Dec 26 22:54:17 2002	P
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Fri Dec 27 12:00:35 2002	P
[1106] CHL:2 ID:72 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 14:56:19 2002	P
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 16:11:32 2002	P
[1106] CHL:2 ID:70 SCSI Target ALERT: SCSI Parity/CRC Error Detected	
Sat Dec 28 17:52:19 2002	P

Note – The controller can store up to 1000 event logs. An event log can record a configuration or operation event as well as an error message or alarm event.

2. Use your arrow keys to move up and down through the list.

- To clear the events from the log once you've read them, use your arrow keys to move down to the last event you want to clear and press Return.

A "Clear Above xx Event Logs?" confirmation message is displayed.

Event Logs		
[1106] CHL:2 ID:74 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
[1]	Clear Above 14 Event Logs ?	ERT: SCSI Parity/CRC Error Detected
[1]	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	ERT: SCSI Parity/CRC Error Detected
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Thu Dec 26 22:54:17 2002		
[1106] CHL:2 ID:0 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Fri Dec 27 12:00:35 2002		
[1106] CHL:2 ID:72 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 14:56:19 2002		
[1106] CHL:2 ID:2 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 16:11:32 2002		
[1106] CHL:2 ID:70 SCSI Target ALERT: SCSI Parity/CRC Error Detected		
Sat Dec 28 17:52:19 2002		

- Choose Yes to clear the recorded event logs.

Note – Resetting the controller clears the recorded event logs. To retain event logs after controller resets, you can install and use the Sun StorEdge Configuration Service program.

7.5 Restoring Your Configuration (NVRAM) From a File

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

The NVRAM configuration file restores all configuration settings (channel settings, and host IDs) but does not rebuild logical drives. See ["Saving Configuration \(NVRAM\) to a Disk" on page 5-59](#) for information about how to save a configuration file, including advice on saving controller-dependent configuration whenever a configuration change is made.

See ["Record of Settings" on page D-1](#) for advice about keeping a written record of your configuration before saving or restoring configuration files. See ["Save NVRAM to Disk and Restore From Disk" on page D-7](#) for a convenient place to keep records whenever you save or restore configuration files.



Caution – Before restoring a configuration file, be certain that the configuration file you apply matches the array to which you apply it. If host IDs, logical drive controller assignments, or other controller-dependent configuration information described in the [Chapter 5](#) has changed since the configuration file was saved, you might lose access to mismatched channels or drives. You have to change cabling or host or drive channel IDs to correct this mismatch and restore the access you have lost. On host Solaris workstations, the address of the RAID controller channel must also match what is described in `/etc/vfstab`.

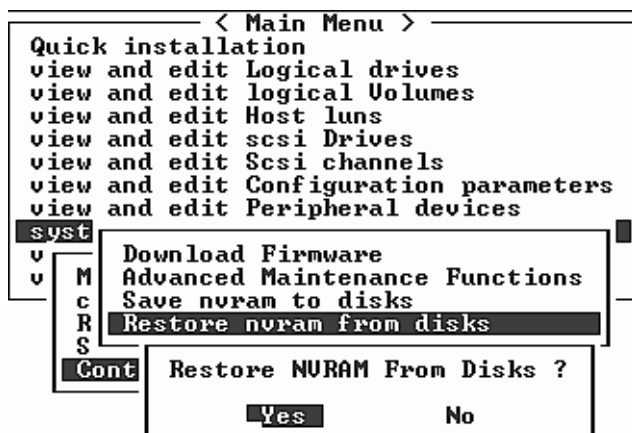
Note – In the Sun StorEdge Configuration Service program, you can save a configuration file that can restore all configurations and rebuild all logical drives. However, it also erases all data when it rebuilds all logical drives, so this operation should be performed only when no data has been stored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

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

1. Choose “system Functions → Controller maintenance → Restore nvram from disks.”

```
      < Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
u
v
  Mute beeper
  change Password
  Reset controller
  Shutdown controller
  Controller maintenance
```

2. Choose Yes to confirm.



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

7.6 Upgrading Firmware

From time to time, firmware upgrades are made available as patches that you can download from SunSolve Online, located at:

<http://sunsolve.sun.com>

Each patch applies to one or more particular pieces of firmware, including:

- Controller firmware
- Drive firmware
- SES firmware (Fibre Channel only)
- PLD firmware (Fibre Channel only)
- SAF-TE firmware (SCSI only)

SunSolve has extensive search capabilities that can help you find these patches, as well as regular patch reports and alerts to let you know when firmware upgrades and other patches become available. In addition, SunSolve provides reports about bugs that have been fixed in patch updates.

Each patch includes an associated README text file that provides detailed instructions about how to download and install that patch. But, generally speaking, all firmware downloads follow the same steps:

- Locating the patch on SunSolve that contains the firmware upgrade you want
- Downloading the patch to a location on your network

- Using your array software (Sun StorEdge Configuration Service or `sscli[1M]` or array firmware, in some cases, to “flash” the firmware to the device it updates

Note – For instructions on how to download firmware to disk drives in a JBOD directly attached to a host, refer to the README file in the patch that contains the firmware.



Caution – Be particularly careful about downloading and installing PLD firmware. If the wrong firmware is installed, or the firmware is installed on the wrong device, your controller might be rendered inoperable. Always be sure to upgrade your SES firmware first before trying to determine if you need a PLD upgrade.

7.6.1 Patch Downloads

1. Once you have determined that a patch is available to update firmware on your array, make note of the patch number or use SunSolve Online’s search capabilities to locate and navigate to the patch.
2. Read the README text file associated with that patch for detailed instructions on downloading and installing the firmware upgrade.
3. Follow those instructions to download and install the patch.

7.6.2 Controller Firmware Upgrade Features

The following firmware upgrade features apply to the controller firmware:

- Redundant Controller Rolling Firmware Upgrade

When downloading is performed on a dual-controller system, firmware is flashed onto both controllers without interrupting host I/O. When the download process is complete, the primary controller resets and lets the secondary controller take over the service temporarily. When the primary controller comes back online, the secondary controller hands over the workload and then resets itself for the new firmware to take effect. The rolling upgrade is automatically performed by controller firmware, and the user’s intervention is not necessary.

- Automatically Synchronized Controller Firmware Versions

A controller that replaces a failed unit in a dual-controller system is often running a newer release of firmware version. To maintain compatibility, the surviving primary controller automatically updates the firmware running on the replacement secondary controller to the firmware version of the primary controller.

Note – When you upgrade your controller firmware, in the Solaris operating environment the `format(1M)` command still shows the earlier revision level. To correct this, you need to update the drive label, using the `autoconfigure` option (option 0) of the `format(1M)` command. When you select `label`, the drive is labelled with the updated firmware version.

- Upgrading Firmware Using the Serial Port Connections (from Microsoft Windows hosts)

The firmware can be downloaded to the RAID controller by using an ANSI/VT100-compatible emulation program. The emulation program must support the ZMODEM file transfer protocol. Emulation programs such as HyperTerminal, Telix, and PROCOMM Plus can perform the firmware upgrade.

7.6.3 Installing Firmware Upgrades

It is important that you run a version of firmware that is supported for your array.



Caution – Before updating your firmware, make sure that the version of firmware you want to use is supported for your array. Refer to the release notes for your array for Sun Microsystems patches containing firmware upgrades that are available for your array. See SunSolve Online for subsequent patches containing firmware upgrades.

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

To download new versions of controller firmware, disk drive firmware, or SES and PLD firmware, use one of the following tools:

- Sun StorEdge CLI (with an in-band connection, for Linux and Microsoft Windows hosts, and for servers running the Solaris operating environment)
- Sun StorEdge Configuration Service program (with an in-band connection, for Solaris and Microsoft Windows hosts)
- Firmware application (only for downloading controller firmware from a Microsoft Windows host with an out-of-band serial port connection)



Caution – You should not use both in-band and out-of-band connections at the same time to manage the array. You might cause conflicts between multiple operations.

7.6.4 Installing Controller Firmware Upgrades From the Firmware Application

You can use a Microsoft Windows terminal emulation session with ZMODEM capabilities to access the firmware application. To upgrade the RAID controller firmware through the serial port and the firmware application, perform the following steps:

1. **Establish the serial port connection.**
 - To upgrade both the boot record and the firmware binary, perform Step 2.
 - To upgrade only the firmware binary, perform Step 3.
2. **Upgrade both boot record and firmware binaries with the following steps:**
 - a. **Choose “system Functions → Controller Maintenance → Advanced Maintenance Functions → Download Boot Record and Firmware.”**
 - b. **Set ZMODEM as the file transfer protocol of your emulation software.**
 - c. **Send the Boot Record Binary to the controller. In HyperTerminal, go to the “Transfer” menu and choose “Send file.”**

If you are not using HyperTerminal, choose “Upload” or “Send” (depending on the software).
 - d. **After the Boot Record Binary has been downloaded, send the Firmware Binary to the controller. In HyperTerminal, go to the “Transfer” menu and choose “Send file.”**

If you are not using HyperTerminal, choose “Upload” or “Send” (depending on the software).

When the firmware update is complete, the controller automatically resets itself.
3. **Upgrade the firmware binary with the following steps:**

- a. Choose “System Functions → Controller maintenance → Download Firmware.”

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit scsi Drives
view and edit Scsi channels
view and edit Configuration parameters
view and edit Peripheral devices
syst
  v
  v
  M Download Firmware
  c Advanced Maintenance Functions
  R Save nvrAm to disks
  S Restore nvrAm from disks
Controller maintenance
```

- b. Set ZMODEM as the file transfer protocol of your emulation software.
- c. Send the firmware binary to the controller. In HyperTerminal, choose “Send file.”

If you are not using HyperTerminal, choose “Upload” or “Send” (depending on the software).

When the firmware update is complete, the controller automatically resets itself.

7.7 Replacing the Front Bezel and Ear Caps

Some procedures require that you remove the front bezel and the small vertical plastic caps on either side of the bezel that cover the rackmount tabs. These rackmount tabs are often referred to as “ears.”

7.7.1 Removing the Front Bezel and Ear Caps

1. Use the provided key to unlock both bezel locks.
2. Grasp the front bezel cover on both sides and pull it forward and then down.

Note – For many operations, including replacing disk drives, it is not necessary to further detach the bezel, since dropping it down moves it sufficiently out of the way.

3. Press the right bezel arm (hinge) toward the left side to release it from the chassis hole.

The left hinge also disengages.

4. Note the location of the chassis bezel holes on each ear.
5. Remove the plastic caps from the front left and right ears of the array.
Both plastic caps are removed in the same way.
 - a. Squeeze both sides of the cap at the top and the bottom.
 - b. Turn the cap toward the center of the array until it disengages and pull it free.

7.7.2 Placing the Bezel and Ear Caps Back Onto the Chassis

Each plastic cap is replaced in the same way, but be sure to place the cap with LED labels on the right ear.

1. Align the inside round notches of the cap with the round cylindrical posts (ball studs) on the ear.
2. Push the top and bottom of the ear cap onto the ear, pressing in on the top side toward the center of the array first.
3. Continue pushing the top and bottom of the ear cap onto the ear, pressing on the side toward the outside of the array.

Do not use force when placing a cap on an ear.

Caution – Be careful to avoid “wedging” the Reset button below the LEDs on the right ear when you replace the plastic cap over it.

4. Insert the bezel arms into the chassis holes.
5. Lift the bezel into position and press it onto the front of the chassis until it is flush with the front.
6. Use the key to lock both bezel locks.

Troubleshooting Your Array

This chapter covers the following maintenance and troubleshooting topics:

- “RAID LUNs Not Visible to the Host” on page 8-1
- “Controller Failover” on page 8-2
- “Rebuilding Logical Drives” on page 8-3
- “Identifying a Failed Drive for Replacement” on page 8-7
- “Recovering From Fatal Drive Failure” on page 8-11
- “SES Temperature Sensor Locations” on page 8-13
- “Modifying Drive-Side SCSI Parameters” on page 8-15
- “Troubleshooting Flowcharts” on page 8-15

For more troubleshooting tips, refer to the Sun StorEdge 3510 FC Release Notes at:

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

8.1 RAID LUNs Not Visible to the Host



Caution – When mapping partitions to LUN IDs, there must be a LUN 0. Otherwise, none of the LUNs will be visible.

By default, all RAID arrays are preconfigured with one or two logical drives. For a logical drive to be visible to the host server, its partitions must be mapped to host LUNs. For mapping details, refer to “[Mapping Logical Drive Partitions to Host LUNs](#)” on page 5-45. Check that you have completed this task.

To make the mapped LUNs visible to a specific host, perform any steps required for your operating system or environment. For host-specific information about different operating environments and operating systems, see:

- “Configuring a Sun Server Running the Solaris Operating Environment” on page F-1
- “Configuring a Windows 2000 Server or Windows 2000 Advanced Server” on page G-1
- “Configuring a Linux Server” on page H-1
- “Configuring an IBM Server Running the AIX Operating Environment” on page I-1
- “Configuring an HP Server Running the HP-UX Operating Environment” on page J-1
- “Configuring a Windows NT Server” on page K-1

8.2 Controller Failover

Controller failure symptoms are as follows:

- The surviving controller sounds an audible alarm.
- The RAID controller status LED on the failed controller is a solid amber.
- The surviving controller sends event messages announcing the controller failure of the other controller.

A “Redundant Controller Failure Detected” alert message is displayed and written to the event log.

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

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

The failover and failback processes are completely transparent to the host.

Controllers are hot-swappable if you are using a redundant configuration, and replacing a failed unit takes only a few minutes. Since the I/O connections are on the controllers, you might experience some unavailability between the times when the failed controller is removed and a new one is installed in its place.

To maintain your redundant controller configuration, replace the failed controller as soon as possible. For details, refer to *Sun StorEdge 3000 Family FRU Installation Guide*.

8.3 Rebuilding Logical Drives

This section describes automatic and manual procedures for rebuilding logical drives.

Note – As disks fail and are replaced, the rebuild process regenerates the data and parity information that was on the failed disk. However, the NVRAM configuration file that was present on the disk is not re-created. After the rebuild process is complete, restore your configuration as described in [“Restoring Your Configuration \(NVRAM\) From a File”](#) on page 7-24.

8.3.1 Automatic Logical Drive Rebuild

Rebuild with Spare. When a member drive in a logical drive fails, the controller first examines whether there is a local spare drive assigned to this logical drive. If yes, it automatically starts to rebuild the data of the failed disk to it.

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

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

To enable this feature, follow these steps:

1. Choose “view and edit Configuration parameters → Drive-side SCSI Parameters → Periodic Auto-Detect Failure Drive Swap Check Time.”

When the “Periodic Auto-Detect Failure Drive Swap Check Time” is enabled (that is, when a check time interval has been selected), the controller detects whether the failed drive has been swapped by checking the failed drive's channel/ID. Once the failed drive has been swapped, the rebuild begins immediately.

Note – This feature requires system resources and can impact performance.

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

For a flowchart of automatic rebuild, see [FIGURE 8-1](#).

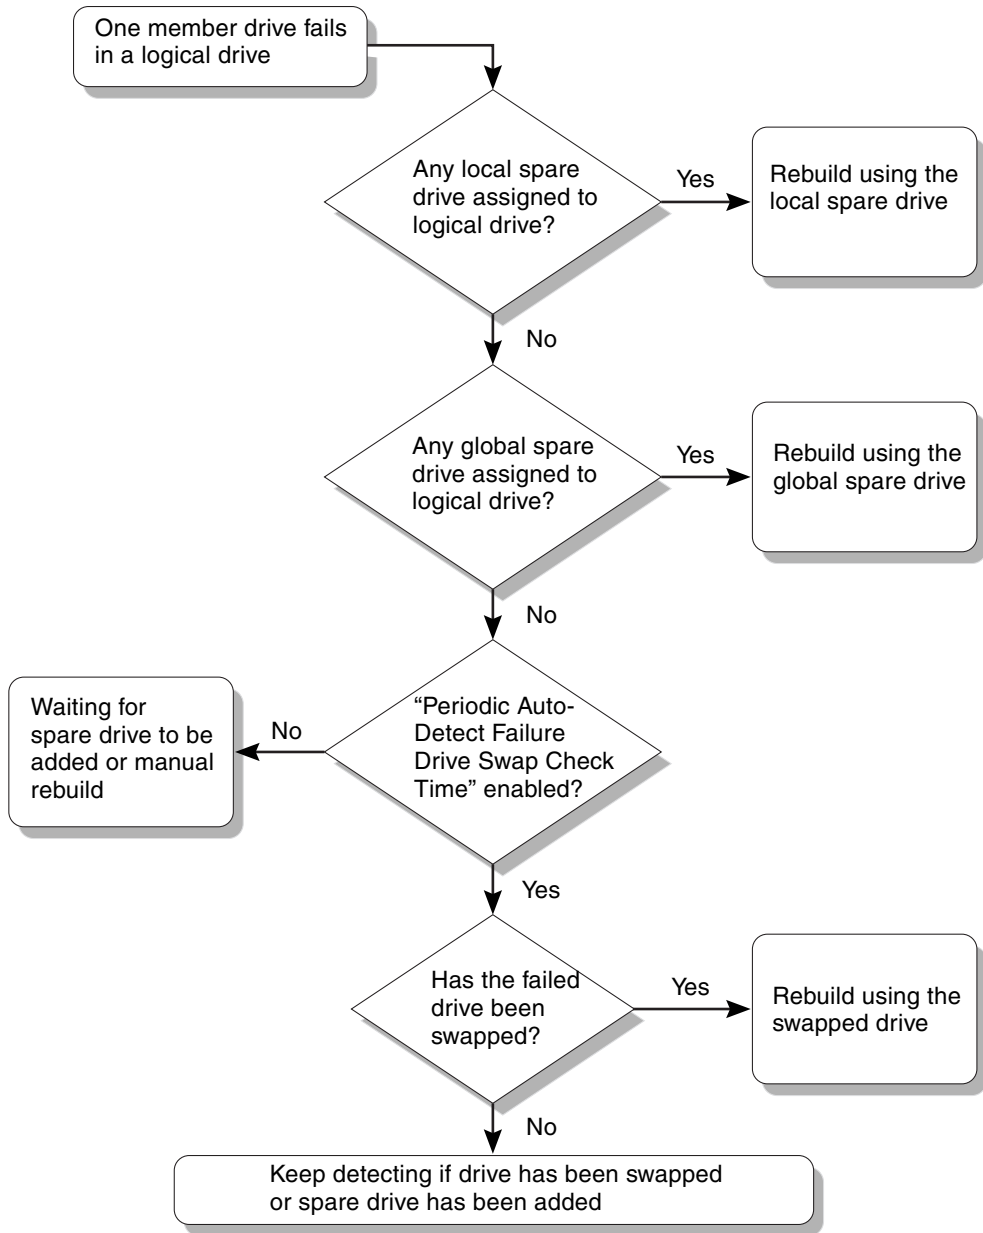


FIGURE 8-1 Automatic Rebuild

8.3.2 Manual Rebuild

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

If there is no local spare available, the controller searches for a global spare. If there is a global spare, the logical drive rebuild begins. See [FIGURE 8-2](#).

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

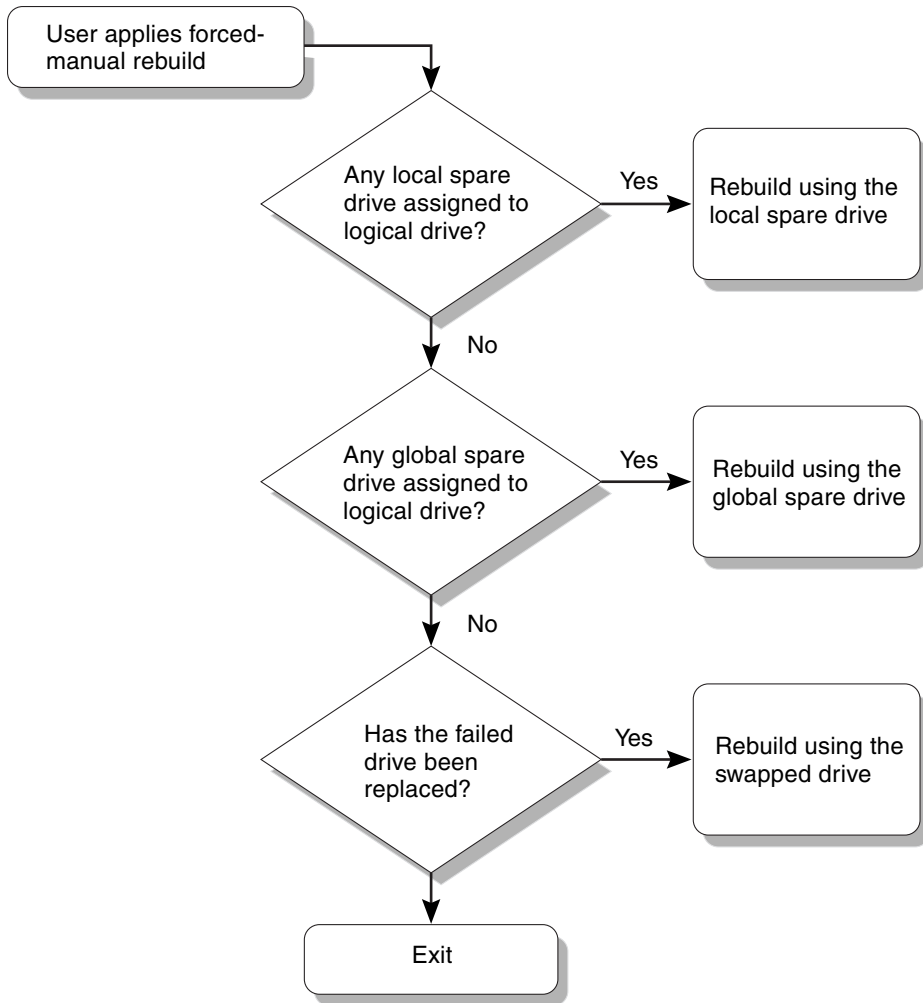


FIGURE 8-2 Manual Rebuild

8.3.3 Concurrent Rebuild in RAID 1+0

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

8.4 Identifying a Failed Drive for Replacement

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



Caution – If when trying to remove a failed drive, you mistakenly remove the wrong drive in the same logical drive, you will no longer be able to access the logical drive. You have incorrectly failed a second drive and caused a critical failure of the RAID set.

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

To find a failed drive, identify a single drive, or test all drive activity LEDs, you can flash the LEDs of any or all drives in an array. Since a defective drive does not light up, this provides a good way to visually identify a failed drive before replacing it.

1. Choose “view and edit scsi Drives.”

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

2. Select the drive you want to identify, and then press Return.

3. Choose "Identify scsi drive → flash All drives" to flash the activity LEDs of all of the drives in the drive channel.

Quic	Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID	
view		2(3)	0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
view				View drive information			0	ON-LINE	SEAGATE ST336752FSUN36G
view				Scan scsi drive			0	ON-LINE	SEAGATE ST336752FSUN36G
view				set slot Number			0	ON-LINE	SEAGATE ST336752FSUN36G
view				add drive Entry			0	ON-LINE	SEAGATE ST336752FSUN36G
view				Identify scsi drive			0	ON-LINE	SEAGATE ST336752FSUN36G
view				flash All drives			ON-LINE	SEAGATE ST336752FSUN36G	
view				flash Selected drive			ON-LINE	SEAGATE ST336752FSUN36G	
view				flash all But selected drive			ON-LINE	SEAGATE ST336752FSUN36G	
		2(3)	6	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	
		2(3)	7	34732	200MB	1	ON-LINE	SEAGATE ST336752FSUN36G	

The option to change the Flash Drive Time is displayed.

Quic	Slot	Chl	ID	Size(MB)	Speed	LG_DRU	Status	Vendor and Product ID	
view		2(3)	0	34732	200MB	0	ON-LINE	SEAGATE ST336752FSUN36G	
view				View drive information			0	ON-LINE	SEAGATE ST336752FSUN36G
view				Scan scsi drive			0	ON-LINE	SEAGATE ST336752FSUN36G
view				set slot Number			0	ON-LINE	SEAGATE ST336752FSUN36G
view				add drive Entry			0	ON-LINE	SEAGATE ST336752FSUN36G
view				Identify scsi drive			0	ON-LINE	SEAGATE ST336752FSUN36G
view				flash All drives			ON-LINE	SEAGATE ST336752FSUN36G	
view				flash Selected drive			E	SEAGATE ST336752FSUN36G	
		2(3)		Flash Drive Time(Second) : 15_			E	SEAGATE ST336752FSUN36G	
		2(3)					E	SEAGATE ST336752FSUN36G	

4. Change the duration if you want. Then press Return and choose Yes.

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

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

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

8.4.1 Flash Selected Drive

When you choose this menu option, the read/write LED of the drive you select flashes for a configurable period of time from 1 to 999 seconds.




FIGURE 8-3 Flashing the Drive LED of a Selected Drive

8.4.2 Flash All SCSI Drives

The “Flash All SCSI Drives” menu option flashes LEDs of all good drives but does not flash LEDs for any defective drives. In the illustration, there are no defective drives.




FIGURE 8-4 Flashing All Drive LEDs to Detect a Defective Non-Flashing Drive

8.4.3 Flash All But Selected Drive

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




FIGURE 8-5 Flashing All Drive LEDs Except a Selected Drive LED

8.5 Recovering From Fatal Drive Failure

With a redundant RAID array system, your system is protected with the RAID parity drive and a global spare or spares.

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

In the extremely rare occurrence that two or more drives appear to fail at the same time, perform the following steps:

1. **Discontinue all I/O activity immediately.**
2. **To cancel the beeping alarm, in the firmware Main Menu, choose “system Functions → Mute beeper.”**
See [Section 7.2, “Silencing Audible Alarms” on page 7-5](#) for more information about silencing audible alarms.
3. **Physically check that all the drives are firmly seated in the array and that none have been partially or completely removed.**
4. **In the firmware Main Menu, choose “view and edit Logical drives,” and look for:**
Status: FAILED DRV (one failed drive)
Status: FATAL FAIL (two or more failed drives)

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

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

6. Do one of the following:

- Choose “system Functions → Reset controller” and choose Yes to confirm.
- Power off the array. Wait five seconds, and power on the array.

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

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

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

- If the status is “FAILED DRV,” manually rebuild the RAID set (refer to [“Manual Rebuild” on page 8-5](#)).
- If the status is still “FATAL FAIL,” you have lost all data on the logical drive and must re-create the logical drive. Proceed with the following procedures:
 - a. **Replace the failed drive. Refer to *Sun StorEdge 3000 Family FRU Installation Guide* for more information.**
 - b. **Delete the logical drive. Refer to the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for more information.**
 - c. **Create a new logical drive. See [“Creating Logical Drives \(Optional\)” on page 5-31](#) for more information.**

Note – As disks fail and are replaced, the rebuild process regenerates the data and parity information that was on the failed disk. However, the NVRAM configuration file that was present on the disk is not re-created. After the rebuild process is complete, restore your configuration as described in [“Restoring Your Configuration \(NVRAM\) From a File” on page 7-24](#).

For additional troubleshooting tips, refer to the *Sun StorEdge 3510 FC Family Release Notes*, located at:

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

8.6 SES Temperature Sensor Locations

Monitoring temperature at different points within the array is one of the most important SES functions. High temperatures can cause significant damage if they go unnoticed. If you are alerted to out-of-limits temperatures by one of the SES sensors, choose “view and edit Peripheral devices → View Peripheral Device Status → SES Device → Temperature Sensors” to get information about the sensor reporting the condition.

There are a number of different sensors at key points within the enclosure. The following table shows the location of each of those sensors. The Element ID corresponds to the identifier shown when you choose “view and edit Peripheral devices → View Peripheral Device Status → SES Device → Temperature Sensors.”

TABLE 8-1 Sun StorEdge 3510 Temperature Sensor Locations

Element ID	Description
0	Drive Midplane Left Temperature Sensor #1
1	Drive Midplane Left Temperature Sensor #2
2	Drive Midplane Center Temperature Sensor #3
3	Drive Midplane Center Temperature Sensor #4
4	Drive Midplane Right Temperature Sensor #5
5	Drive Midplane Right Temperature Sensor #6
6	Upper IOM Left Temperature Sensor #7
7	Upper IOM Left Temperature Sensor #8
8	Lower IOM Temperature Sensor #9
9	Lower IOM Temperature Sensor #10
10	Left PSU Temperature Sensor #11
11	Right PSU Temperature Sensor #12

8.7 Identifying Fans From the SES Device Status Menu

Using controller firmware menu options, you can view the status of SES components, including the pair of fans located in each fan and power supply module. A fan is identified by the SES Device menus as a cooling element.

To view the status of each fan, perform the following steps:

1. Choose “view and edit Peripheral devices → View Peripheral Device Status → SES Device → Cooling element.”
2. Select one of the elements (element 0, 1, 2, or 3) and press Return.

Standard fan speeds are indicated by numbers 1 through 7, indicating speeds in the normal range of 4000 to 6000 RPM. The number 0 indicates that the fan has stopped.

If a fan fails and the Status field does not display the OK value, you must replace the fan and power supply module.

Cooling elements in the status table can be identified for replacement as shown in [TABLE 8-2](#).

TABLE 8-2 Relationship Between Cooling Elements, Fans, and Power Supply Modules

Cooling Element #	Fan # and Power Supply Module #
Cooling Element 0	FAN 0, PS 0
Cooling Element 1	FAN 1, PS 0
Cooling Element 2	FAN 2, PS 1
Cooling Element 3	FAN 3, PS 1




FIGURE 8-6 Cooling Fan Locations

8.8 Modifying Drive-Side SCSI Parameters

There are a number of interrelated drive-side SCSI parameters you can set using the “view and edit Configuration parameters” menu option. It is possible to encounter undesirable results if you experiment with these parameters, so it is good practice to only change parameters when you have good reason to do so. Refer to the “Viewing and Editing Configuration Parameters” chapter of the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for cautions about particular parameter settings that should be avoided. In particular, do not set the “Periodic SAF-TE and SES Device Check Time” to less than one second, and do not set the “SCSI I/O Timeout” to anything less than 15 seconds, and preferably no less than the FC default of 30 seconds.

8.9 Troubleshooting Flowcharts

This section provides troubleshooting flowcharts to illustrate common troubleshooting methods.

The flowcharts included in this section are:

- [“Power Supply and Fan Module” on page 8-16](#)
- [“Drive LEDs” on page 8-19](#)
- [“Front Panel LEDs” on page 8-21](#)
- [“I/O Controller Module” on page 8-25](#)

For the JBOD and expansion unit flowchart, refer to [“Troubleshooting Sun StorEdge 3510 FC JBOD Arrays”](#) on page B-22.

For overview information about LEDs, see [Chapter 6](#).

For more information about replacing each module, refer to the *Sun StorEdge 3000 Family FRU Installation Guide for 2U Arrays*.

8.9.1 Power Supply and Fan Module

The following flowchart provides troubleshooting procedures for the power supply and fan module.

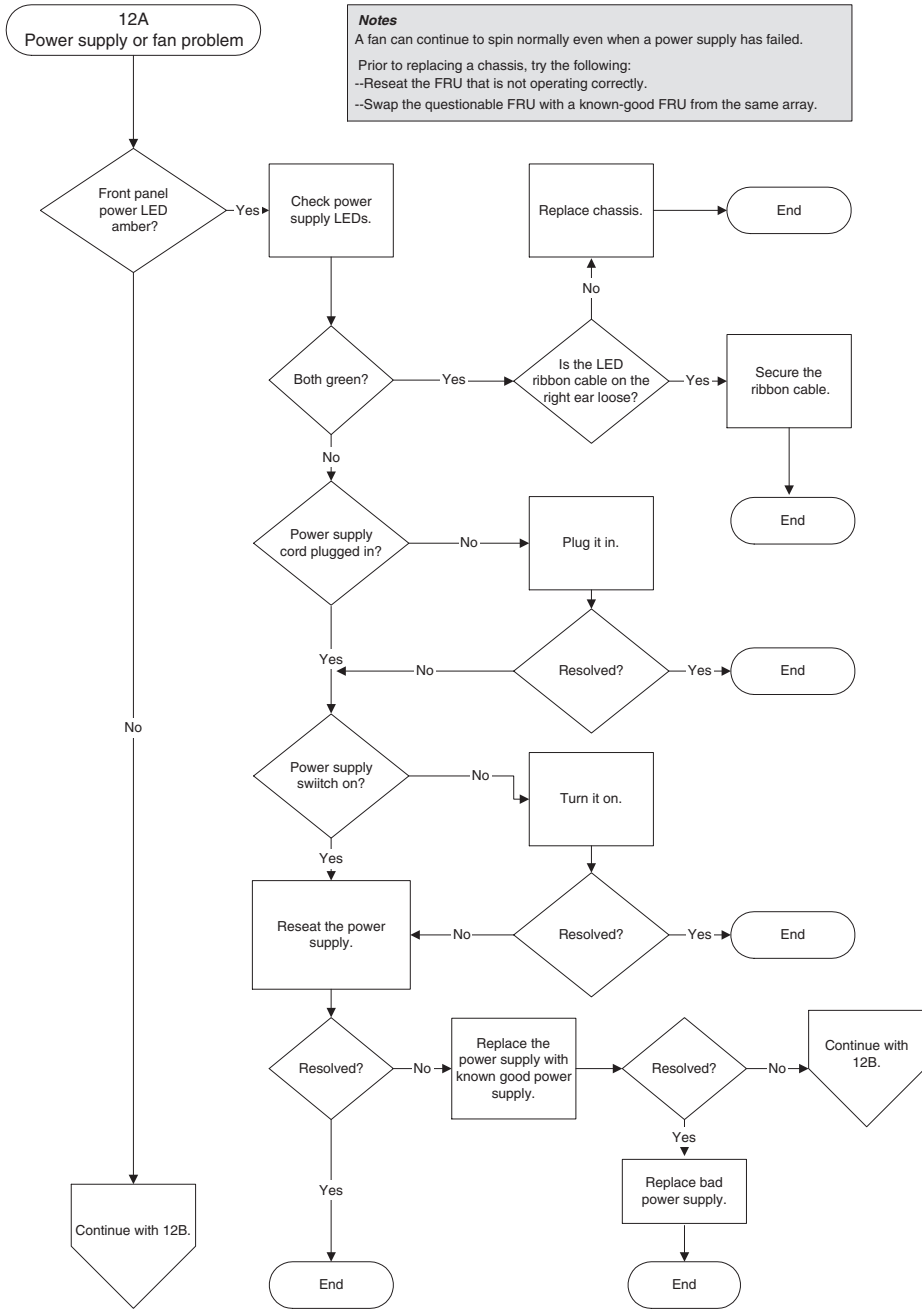


FIGURE 8-1 Power Supply or Fan Module Flowchart, 1 of 2

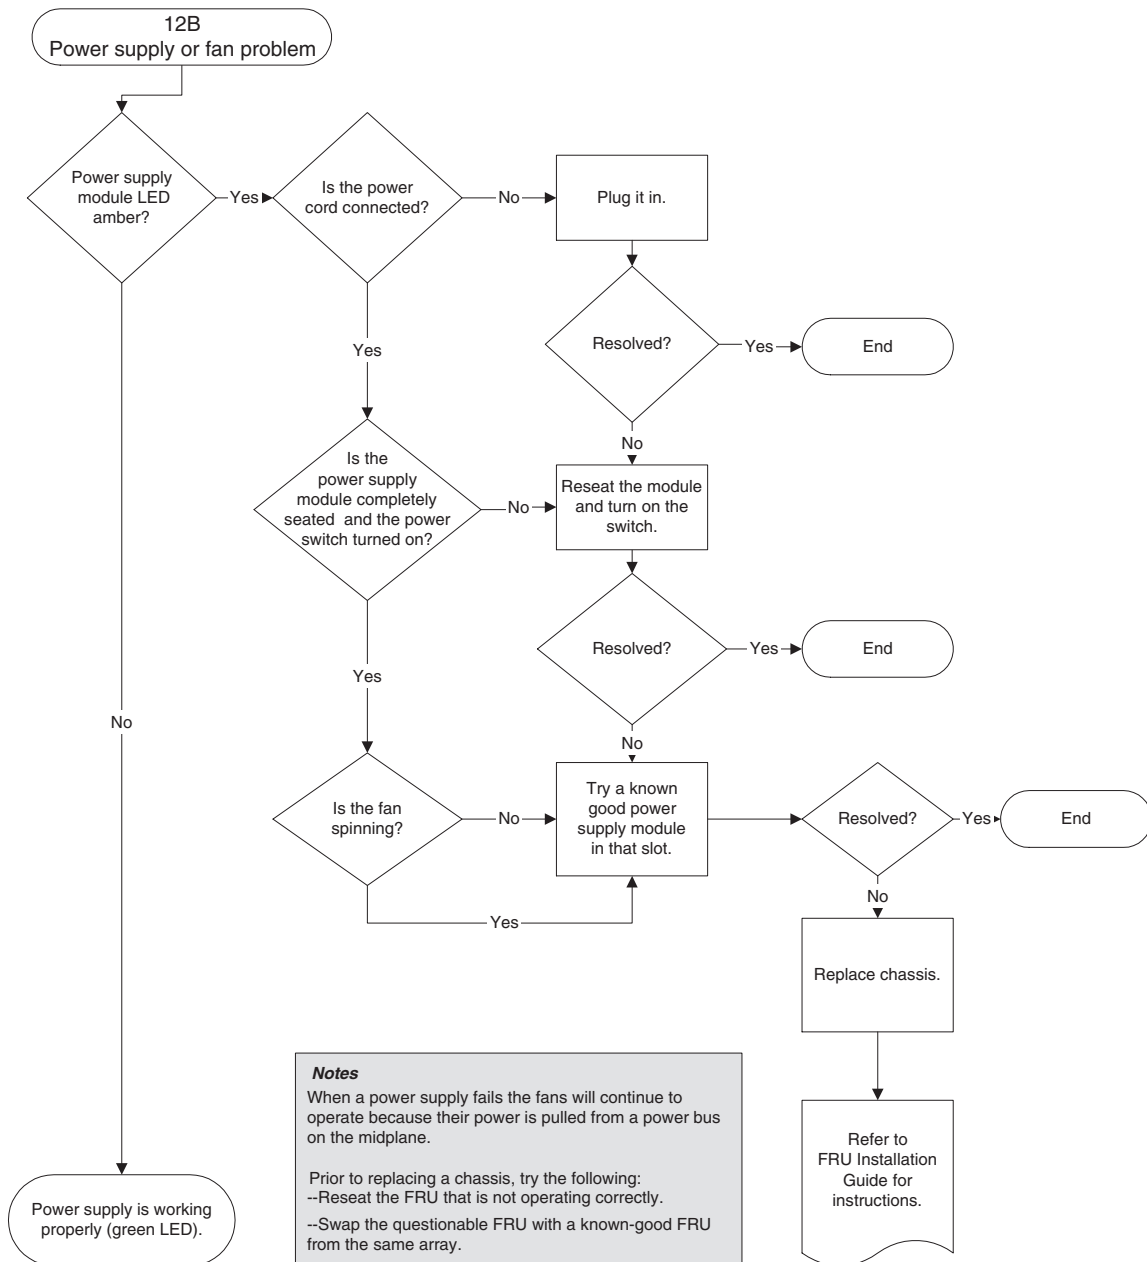


FIGURE 8-2 Power Supply or Fan Module Flowchart, 2 of 2

8.9.2 Drive LEDs

Before you perform the drive LED troubleshooting procedures, you might want to use the firmware application to identify a failed drive. See [“Identifying a Failed Drive for Replacement” on page 8-7](#) for more details.

For overview information about drive LEDs and how they work, see [“Front Panel LEDs” on page 6-1](#).

You can check the physical drive parameters using the firmware application. From the firmware Main Menu, choose “view and edit scsi drives.” For more information about the firmware application, refer to the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for your array.



Caution – When you rotate or replace drives, make sure that:

- All I/O is stopped.
- The “Periodic Drive Check Time” setting in the firmware application is set to disabled (this is the default setting). This prevents automatic drive rebuild, which is not recommended for live systems or troubleshooting.

For more information, refer to “Periodic Drive Check Time” in the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for your array.



Caution – To prevent any possible data loss, back up the chassis data onto another storage device prior to replacing a disk drive.

The following flowchart provides troubleshooting procedures for the FC drive LEDs.

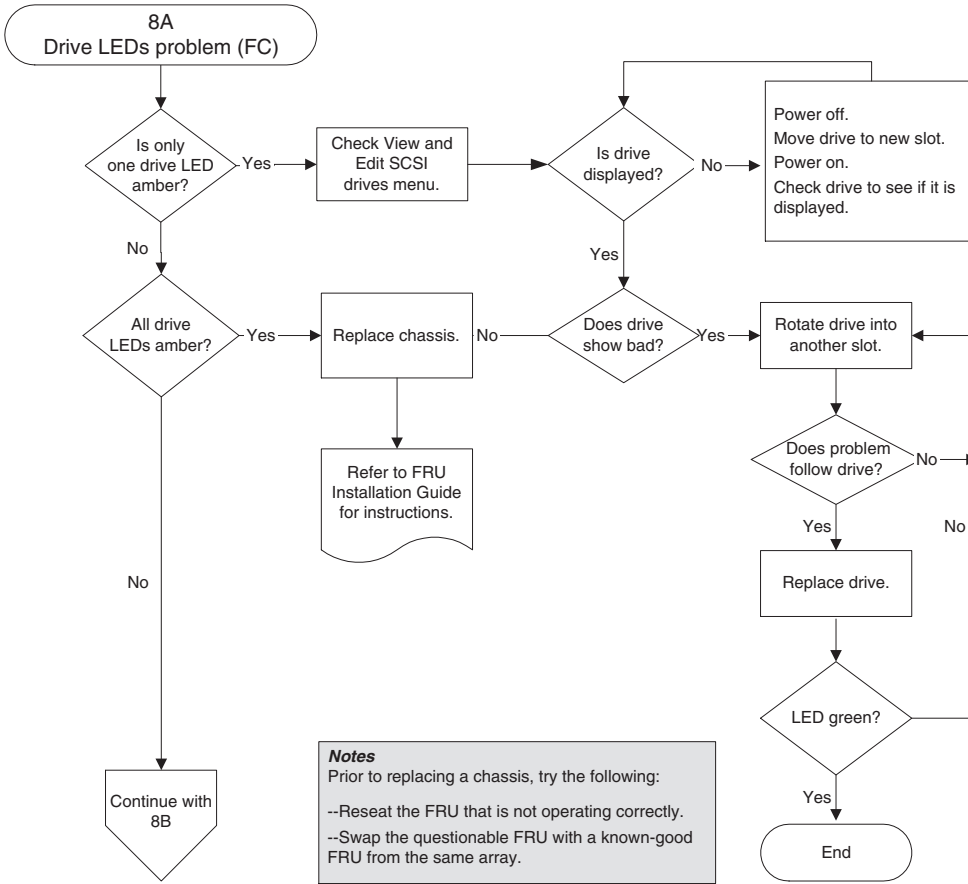


FIGURE 8-3 FC Drive LEDs Flowchart, 1 of 2

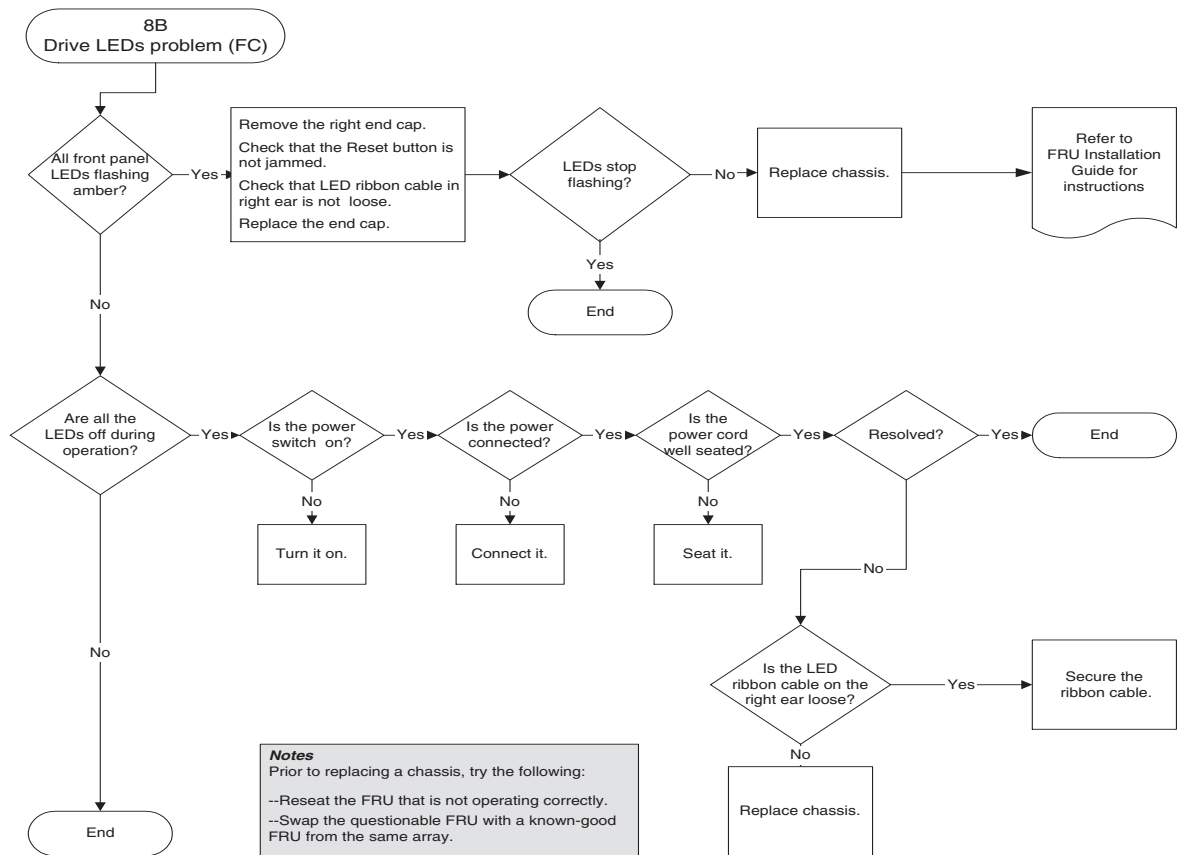


FIGURE 8-4 FC Drive LEDs Flowchart, 2 of 2

For more information about replacing the chassis and EMU module, refer to the *Sun StorEdge 3000 Family FRU Installation Guide for 2U Arrays*.

8.9.3 Front Panel LEDs

The following flowchart provides troubleshooting procedures for the FC front panel LEDs.

Note – The LED ribbon cable referred to in this flowchart is the white cable that connects the front panel LEDs to the midplane. It is located on the right front panel ear and is directly attached to the LEDs.

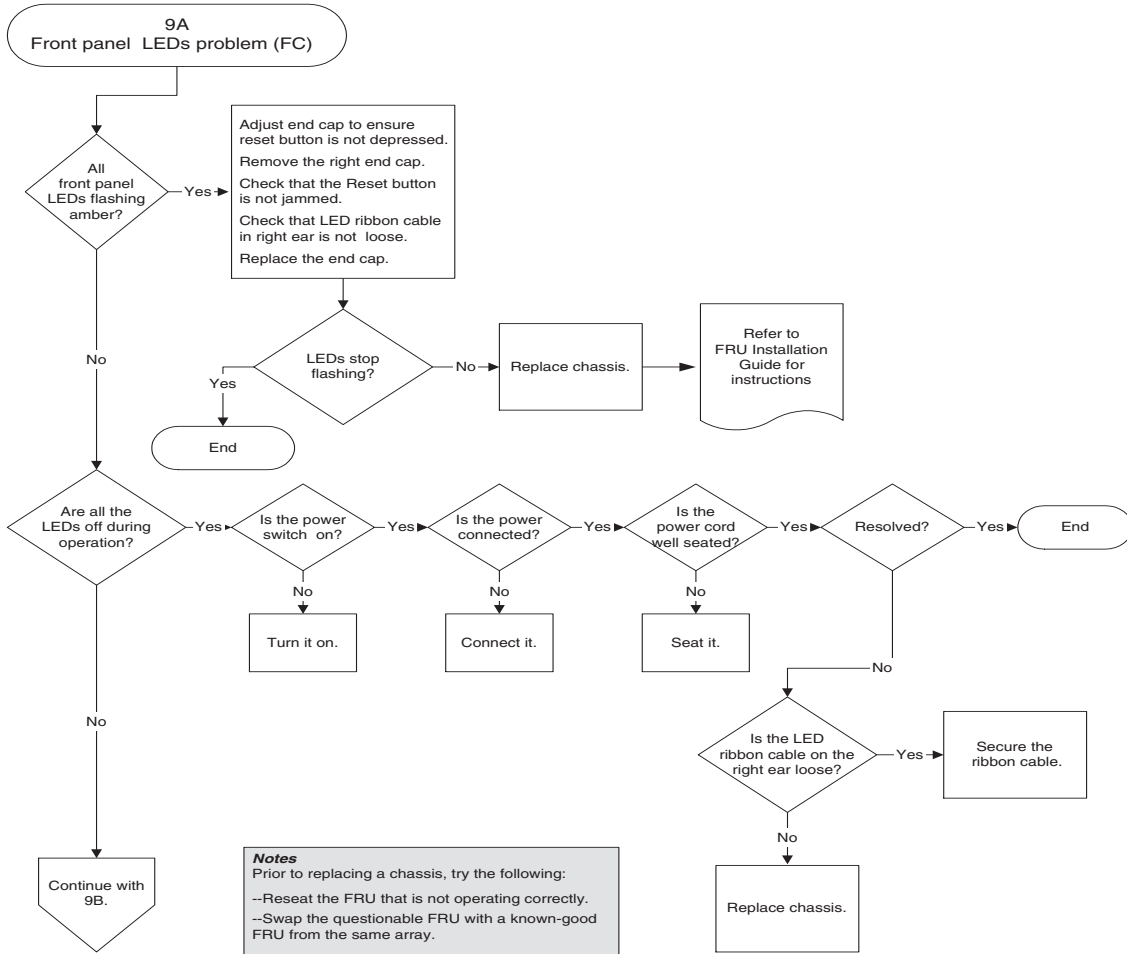


FIGURE 8-5 Front Panel LEDs (FC) Flowchart, 1 of 4

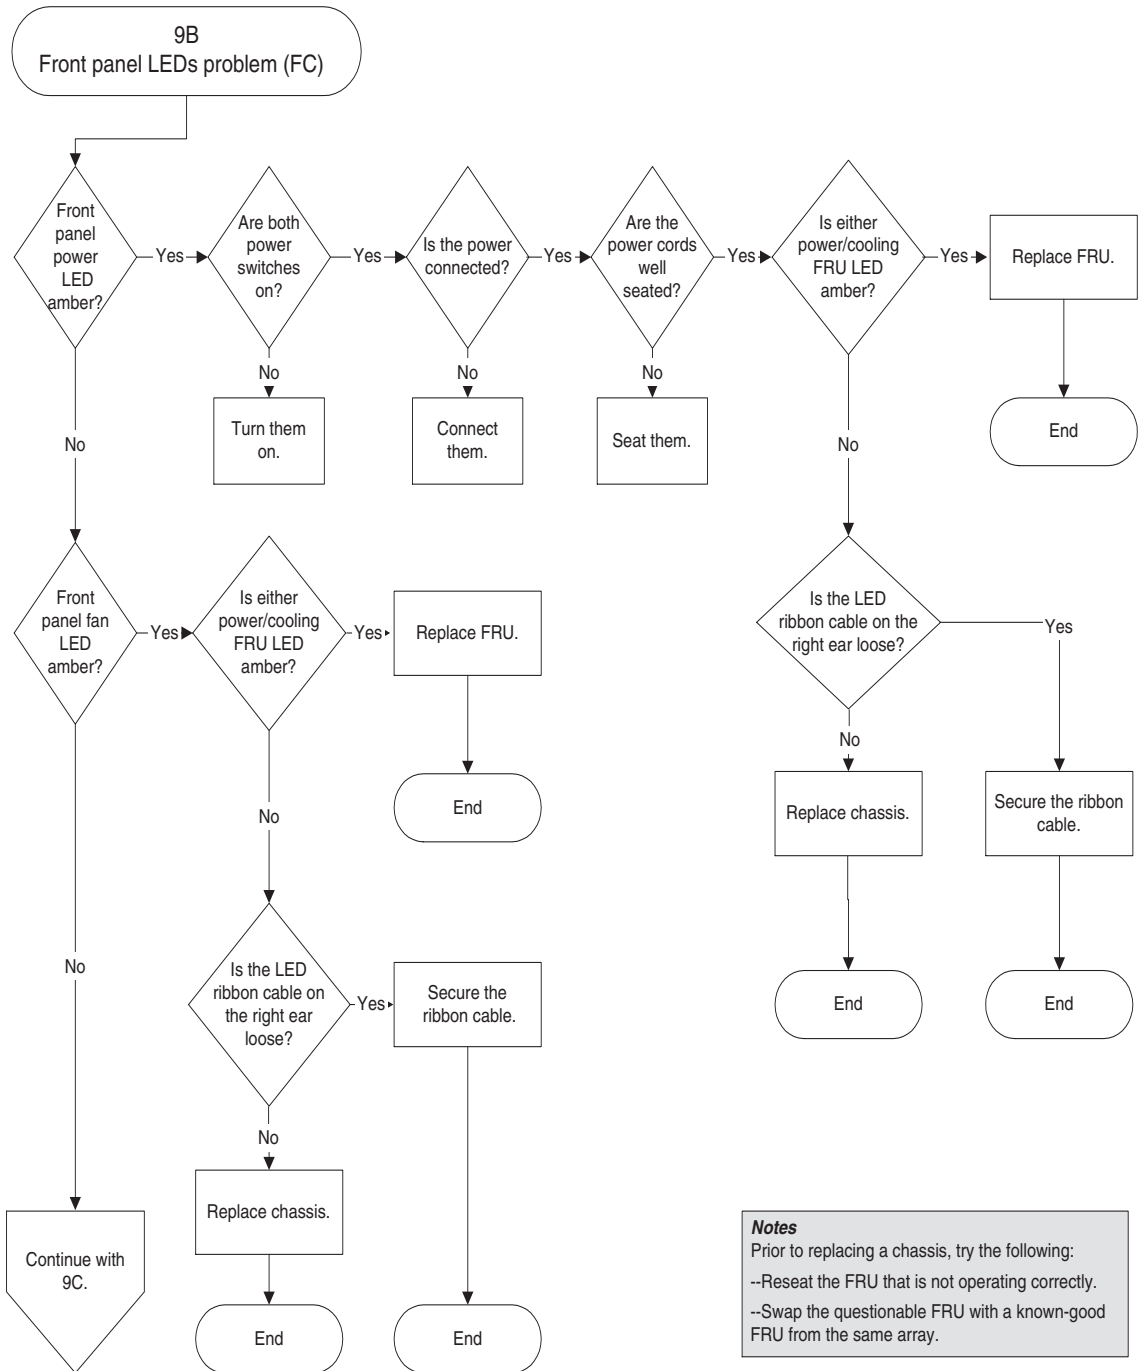


FIGURE 8-6 Front Panel LEDs Flowchart, 2 of 4

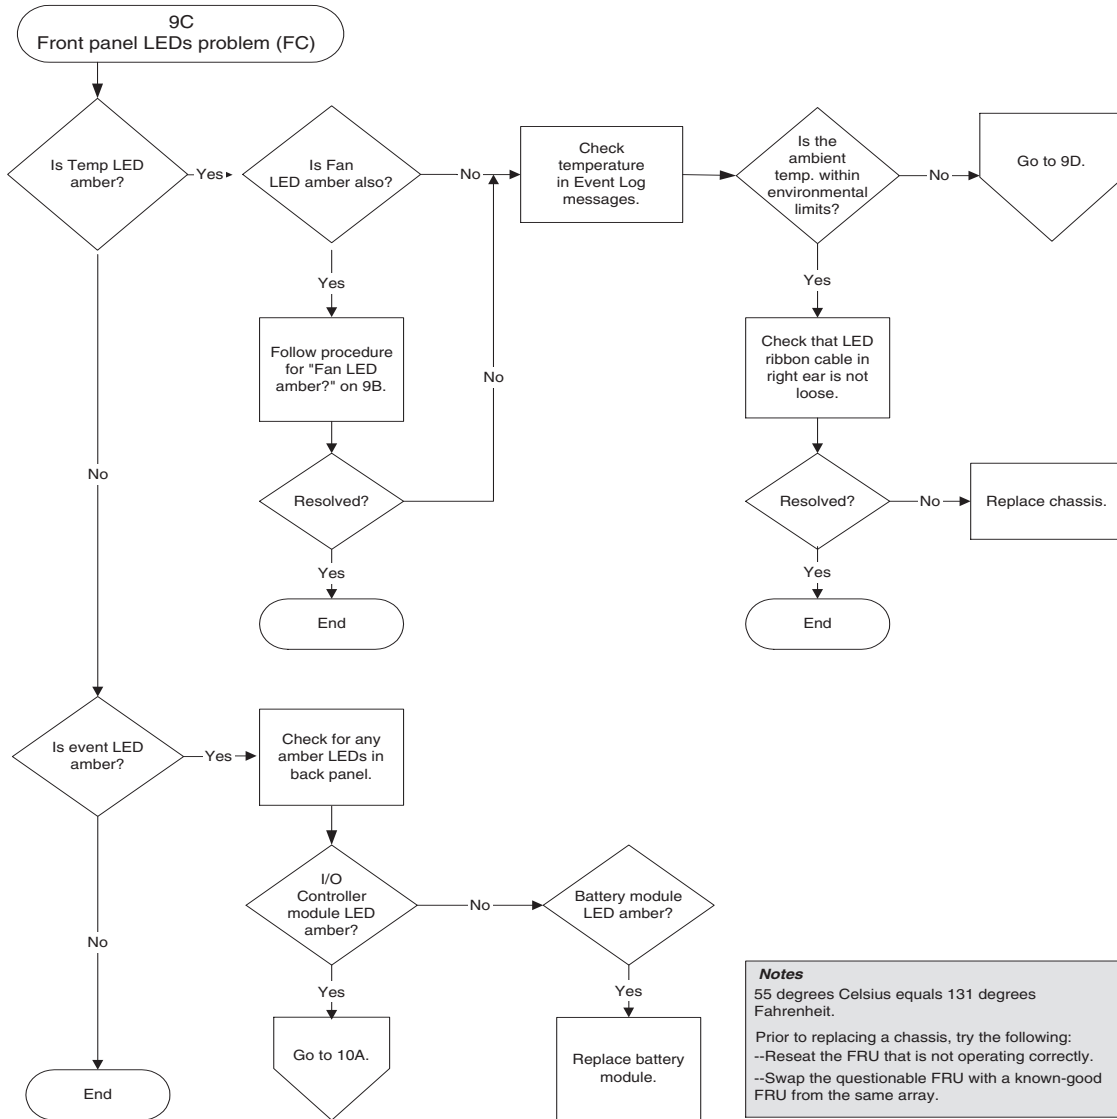


FIGURE 8-7 Front Panel LEDs Flowchart, 3 of 4

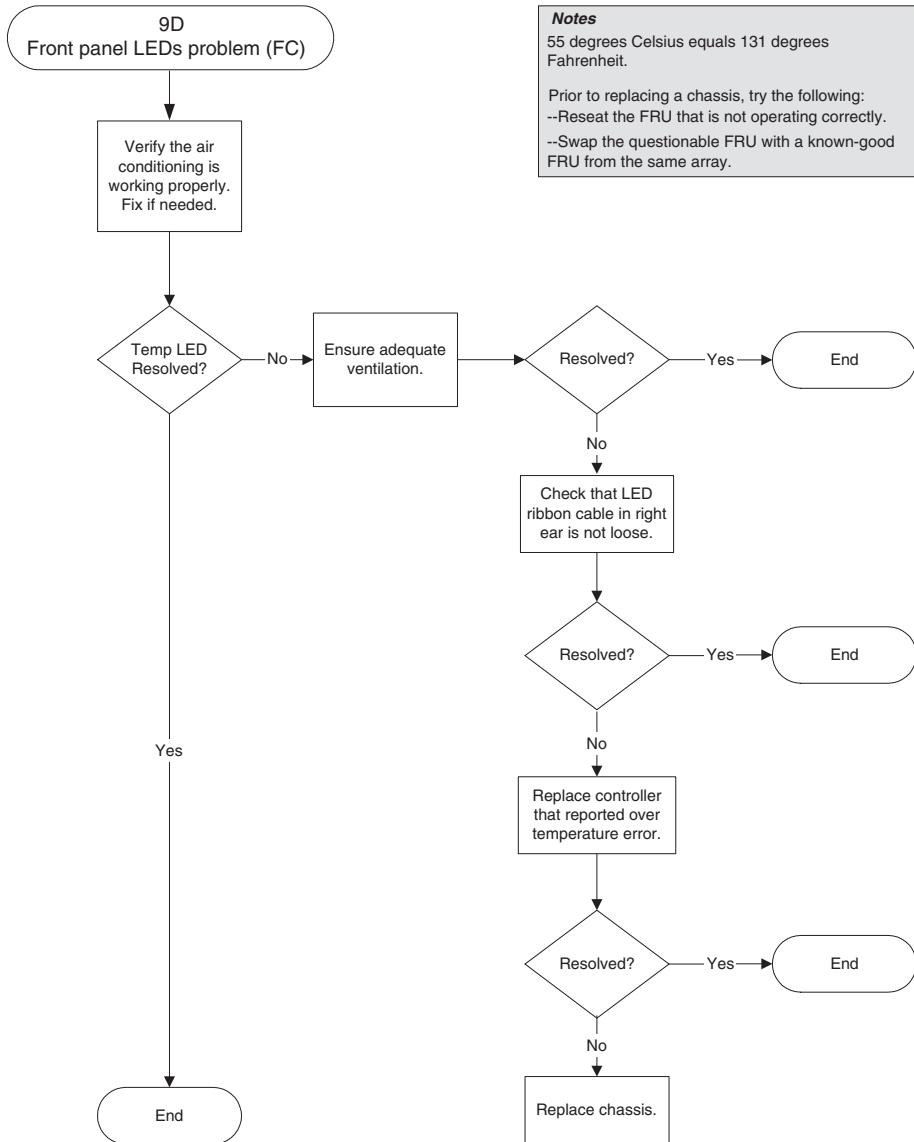


FIGURE 8-8 Front Panel LEDs Flowchart, 4 of 4

8.9.4 I/O Controller Module

The following flowchart provides troubleshooting procedures for the I/O controller module.

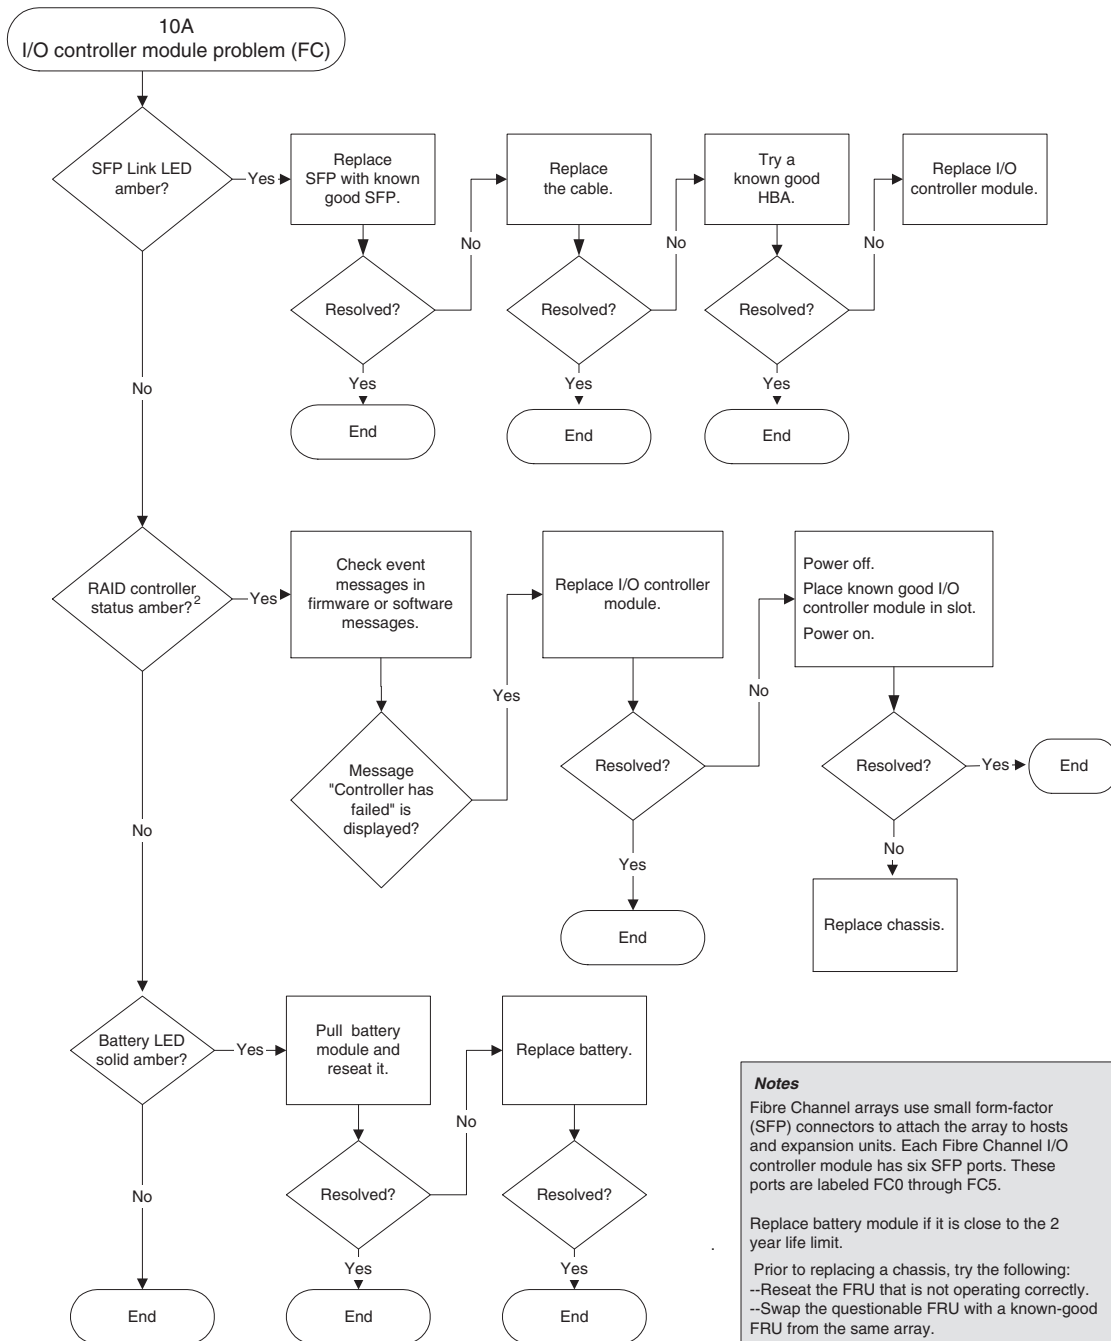


FIGURE 8-9 I/O Controller Module Flowchart

8.10 Using the Reset Button

To test that the LEDs work, using a paperclip, press and hold the Reset button for 5 seconds. All the LEDs should change from green to amber when you perform this test. Any LED that fails to light indicates a problem with the LED. When you release the Reset button, the LEDs return to their initial state. See [“Chassis Ear LEDs and Reset Button on Front Panel” on page 6-2](#) for more information.

To silence audible alarms that are caused by component failures, use a paperclip to push the Reset button. See [Section 7.2, “Silencing Audible Alarms” on page 7-5](#) for more information about silencing audible alarms.

8.11 Silencing Audible Alarms

An audible alarm indicates that either a component in the array has failed or a specific controller event has occurred. The cause of the alarm determines how you silence the alarm. See [Section 7.2, “Silencing Audible Alarms” on page 7-5](#) for more information about silencing audible alarm.

Sun StorEdge 3510 FC Array Specifications

This appendix provides the specifications for the Sun StorEdge 3510 FC array. The topics covered are:

- [“Physical Specifications” on page A-2](#)
- [“Sun StorEdge 3510 FC Array Highlights” on page A-3](#)
- [“Agency Approvals and Standards” on page A-5](#)

A.1 Physical Specifications

TABLE A-1 Sun StorEdge 3510 FC Array Physical Specifications

Description	Desktop	Rackmount
Height	3.64 in. / 9.25 cm	2U (3.45 in. / 8.76 cm)
Width	19 in. / 48.26 cm	17.56 in. / 44.60 cm (body)
Depth	Main chassis 18 in./ 45.72 cm	Main chassis 18 in./ 45.72 cm
	To back of power supply: 20 in. / 50.8 cm	To back of power supply: 20 in. / 50.8 cm
	To back of power supply handle: 21 in. / 53.34 cm	To back of power supply handle: 21 in. / 53.34 cm
Weight (fully loaded)	69.3 lb / 31.4 kg with 73-GB drives	63.3 lb / 28.7 kg with 73-GB drives
Weight (fully loaded expansion unit)	62.9 lb / 28.6 kg with 73-GB drives	56.9 lb / 25.8 kg with 73-GB drives

Note – Add 10.6 lb (4.8 kg) for packaging if you want to know the shipping weight of an array or expansion unit.

A.2 Sun StorEdge 3510 FC Array Highlights

Feature	Description
General	<ul style="list-style-type: none">• Up to 12 hot-pluggable drives in a 2U (3.45 in. / 8.76 cm high) chassis• Autosensing AC or DC power supplies• Dual-host access in certain configurations
Density	<ul style="list-style-type: none">• Up to 876 GB in a RAID array (with 73 GB drives)• 21 in. chassis depth• RAID supports up to two expansion units
Reliability	<ul style="list-style-type: none">• Redundant hot-swappable FRUs• Single or redundant active/active hot-swappable RAID controllers• N+1 hot-swappable power and cooling• NEBS Level 3 and HALT test certified; designed to meet 99.999% reliability
RAID System Storage Resource Management	<ul style="list-style-type: none">• Web-based 10/100BASE-T Ethernet support• Serial port out-of-band support with modem control• RAID Levels 0, 1, 1+0, 3, 3+0, 5, and 5+0• Up to 1024 LUNs• SNMP traps and MIB in firmware• 1-GB cache memory standard• Four channels are user-configurable as drive or host channels; two channels are dedicated drive channels• Support for Sun StorEdge Configuration Service and Diagnostic Reporter management software

The following topics briefly describe specific hardware and firmware features.

A.2.1 Hardware Loop ID

Each device communicating on an arbitrated loop requires an AL_PA, the Arbitrated Loop Physical Address. The host Fibre Channel supports a single hard loop ID in the range from 0 to 125. The drive channel allows up to 125 devices to be connected to a single loop.

The controller supports automatic loop ID assignment on the drive channel. A hard loop address ID is assigned to the disk drives by setting the radial switch.

A.2.2 Firmware Host-Side Connection Mode

The controller default is “Loop Only” in “View and Edit Host-side Parameters.”

A.2.3 Firmware LUN Filtering (RAID-Based Mapping)

LUN filtering is a method used for separating and controlling access to data from the RAID controller. One major benefit of Fibre Channel is the capability to share a common storage pool with multiple servers or workstations. However, allocation becomes an issue when every server in this network can access the data in a disk array. LUN filtering provides a means for controlling data access if data is not allowed for every server.

When you map a logical drive to a host LUN through the firmware application, you can create a mapping according to the WWPN port names that are specific to each host adapter. A LUN mask can then be created as an access filter for including or excluding host adapters from addressing specific storage unit. It is composed of an ID range that is configured to include one or more IDs. A host adapter can be identified by its ID (device-specific port names) as included or excluded from range. The LUN mask is also defined with a filter type as “read only,” “read/write,” or “exclude.”

A.2.4 Firmware Redundant Loops

The selection of “drive-side dual-loop” is fixed. Since two channels are used to connect a group of physical drives, the two channels automatically assume the same ID address and operate in duplex mode.

A.2.5 Firmware Dynamic Load Balancing

The controller provides a mechanism to equally distribute I/O loads between channels in a redundant loop configuration.

A.3 Agency Approvals and Standards

Product Safety	Standard
Country	
U.S.	UL Listed to UL60950:2000, 3rd Edition
Canada	CSA Standard CAN/CSA-C22.2 No. 60950-00 3rd Edition
Germany	TÜV
European Union	EN 60950:2000
Japan	Part of World-wide CB Scheme
Australia	Part of World-wide CB Scheme
Argentina	Resolution 92-98 (S-Mark)
Germany	GS mark (ergonomics) (Rheinland)
Russia	Part of World-wide CB Scheme (GOST-R mark)
Russia	Hygienic Mark (ergonomics)
Electromagnetic Compatibility	
Country	
U.S.	FCC #47, Part 15, Subpart B, Class B
Canada	ICES-003
Japan	VCCI Class B
European Union	EN 55022:1998 Class B
Australia/New Zealand	AS/NZS 3548:1996
Taiwan	BSMI CNS 13438 Class B
Russia	GOST-R mark
Germany	Same as European Union
Argentina	S mark
Electromagnetic Interference	
Test	
Harmonics Emissions	EN 61000-3-2:2000 (No Limits)
Voltage Flicker	EN 61000-3-3:1995/A1:2001 (No Limits)

Product Safety	Standard
ESD	EN 55024 (8kV Contact, 15kV Air)
RF Field	EN 55024 (10V/m)
Electrical Fast Transient Burst	EN 55024 (1kV I/O, 2kV Power)
Surge	EN 55024 (1kV I/O, 1kV Power L-L, 2kV Power L-G)
RF Conducted	EN 55024 (3V I/O and Power)
Power Frequency Magnetic Field	EN 55024 (N/A monitors only)
Voltage Dip and Interruption	EN 55024 (0V/0.5cycle, 70%V/0.5sec, 0V/5sec)

Using Standalone JBOD Arrays

A single Sun StorEdge 3510 FC JBOD array can be connected directly to certain Solaris operating system hosts. The supported configurations differ substantially from the use of Sun StorEdge 3510 FC expansion units connected to Sun StorEdge 3510 FC RAID arrays. This appendix describes the usage of Sun StorEdge 3510 FC JBOD arrays.

The topics covered in this appendix are:

- [“Supported Configurations \(JBOD Arrays\)” on page B-2](#)
- [“Supported Operating Systems \(JBOD Arrays\)” on page B-2](#)
- [“Supported Host Platforms and Connection Methods \(JBOD Arrays\)” on page B-2](#)
- [“Known Limitations Affecting Sun StorEdge 3510 FC JBOD Arrays” on page B-3](#)
- [“Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays” on page B-4](#)
- [“Setting Loop IDs on Sun StorEdge 3510 FC JBOD Arrays” on page B-7](#)
- [“Connecting Sun StorEdge 3510 FC JBOD Arrays” on page B-9](#)
- [“Enabling Sun StorEdge 3510 FC JBOD Array Support” on page B-13](#)
- [“Replacing a Disk Drive in a Sun StorEdge 3510 FC JBOD Array” on page B-15](#)
- [“Understanding Component Status LEDs on JBOD Arrays” on page B-17](#)
- [“Troubleshooting Sun StorEdge 3510 FC JBOD Arrays” on page B-22](#)
- [“Converting JBODs to RAID arrays” on page B-27](#)

B.1 Supported Configurations (JBOD Arrays)

Sun StorEdge 3510 FC JBOD arrays can be attached directly to a single Sun host computer using either single-port or dual-port configurations. Refer to [“Single-Port Connections to Sun StorEdge 3510 FC JBOD Arrays” on page B-9](#) and [“Dual-Port Connections to Sun StorEdge 3510 FC JBOD Arrays” on page B-10](#) for more information.

B.2 Supported Operating Systems (JBOD Arrays)

The Sun StorEdge 3510 FC JBOD array can be used with currently released versions of the Solaris 8 operating system and the Solaris 9 operating system.

B.3 Supported Host Platforms and Connection Methods (JBOD Arrays)

[TABLE B-1](#) lists the Sun systems that are supported for the Sun StorEdge 3510 FC JBOD array. The table also shows the connection methods supported for each of the supported Sun systems.

Supported HBAs include (X)6767A 2 Gbit PCI Single FC HBA and (X)6768A 2 Gbit PCI Dual FC HBA.

TABLE B-1 Supported Sun Servers and Connection Methods (JBOD Arrays)

Server	HBA
Sun Enterprise™ 220R server	Yes
Sun Enterprise 250 server	Yes
Sun Enterprise 420R server	Yes

TABLE B-1 Supported Sun Servers and Connection Methods (JBOD Arrays)

Server	HBA
Sun Enterprise 450 server	Yes
Sun Fire™ V120 server	Yes
Sun Fire V280 server	Yes
Sun Fire V880 server	Yes

B.4 Known Limitations Affecting Sun StorEdge 3510 FC JBOD Arrays

Limitations affecting the Sun StorEdge 3510 FC Array are listed below:

- Booting from a JBOD disk is not supported in this release.
- The `mpxio` functionality of Sun StorEdge SAN Foundation Suite 4.2 software is not supported with Sun StorEdge 3510 FC JBOD arrays.
- Only Sun hosts running the Solaris operating system are supported in this release.
- Sun Cluster software is not supported in this release.
- VERITAS Cluster Server (VCS) software is not supported in this release.
- Only direct-attached single-host connections to a single JBOD are supported in this release. There is no multihost support in this release. Fibre Channel switches are not supported in this release. Daisy-chaining is not supported in this release.
- Sun Storage Automated Diagnostic Environment (StorADE) 2.2 software is not supported in this release.
- Do not use the `luxadm` utility to diagnose or manage Sun StorEdge 3510 FC JBOD arrays.
- Sun StorEdge Configuration Service software supports Sun StorEdge 3510 FC JBOD arrays. However, since Sun StorEdge 3510 FC JBOD arrays do not have a RAID controller or RAID controller firmware to manage the disks, this software support is necessarily limited. Monitoring functionality that does not require a RAID controller or RAID controller firmware works in the same way it does for RAID arrays. See [“Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays” on page B-4](#) for more specific information.
- Similarly, Sun StorEdge Diagnostic Reporter support for standalone JBOD arrays is limited to functionality that can work without a RAID controller or RAID controller firmware. See [“Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays” on page B-4](#) for more specific information.

- The Sun StorEdge CLI supports JBOD arrays. However, because JBOD arrays do not have a RAID controller or RAID controller firmware to manage the disks, this CLI support is limited. See [“Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays”](#) on page B-4 for more specific information.

Refer to the “Monitoring JBODs” appendix in the *Sun StorEdge 3000 Family CLI User’s Guide* for information about using these commands.

B.5 Using Sun StorEdge 3000 Family Software Monitoring and Management Tools with JBOD Arrays

The following software management tools are provided on the Sun StorEdge 3000 Family Professional Storage Manager CD, provided with your array. The Sun StorEdge 3000 Family Documentation CD provides the related user guides.

- **Sun StorEdge Configuration Service.** Provides centralized storage configuration, maintenance, and monitoring functions. Refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide* for in-band setup procedures.
- **Sun StorEdge Diagnostic Reporter.** Provides event monitoring and notification. Refer to the *Sun StorEdge Diagnostic Reporter User’s Guide* for information.
- **Sun StorEdge Command-Line Interface (CLI).** A command-line interface utility that provides script-based management. Refer to the *Sun StorEdge CLI User’s Guide* for CLI information.

For details on how to install Sun StorEdge Configuration Service, Sun StorEdge Diagnostic Reporter, or the Sun StorEdge CLI software, refer to the *Sun StorEdge 3000 Family Software Installation Manual*.

B.5.1 Sun StorEdge Configuration Service

Sun StorEdge Configuration Service supports the Sun StorEdge 3310 SCSI array and Sun StorEdge 3510 FC array. It also supports, to a limited degree, standalone JBOD arrays. Since standalone JBOD arrays do not have a RAID controller or RAID controller firmware to manage the disks, this software support for JBODs is limited to the following functions:

- viewing component and alarm characteristics
- upgrading firmware on hard drives
- upgrading firmware on SAF-TE devices

Refer to the “Monitoring JBODs” appendix in the *Sun StorEdge 3000 Family Configuration Service User’s Guide* for information about using these functions with JBOD arrays.

B.5.1.1 Enabling JBOD Support

Use JBOD support only when you have a JBOD array connected directly to a host. This enables you to monitor peripheral device condition and events.

Note – Enabling JBOD support may impact I/O performance.

To monitor peripheral device condition and events of a Sun StorEdge FC JBOD array or Sun StorEdge SCSI JBOD array from the SCS Console, you first need to enable JBOD support.

1. Select View → Agent Options Management.

The Agent Options Management window is displayed.

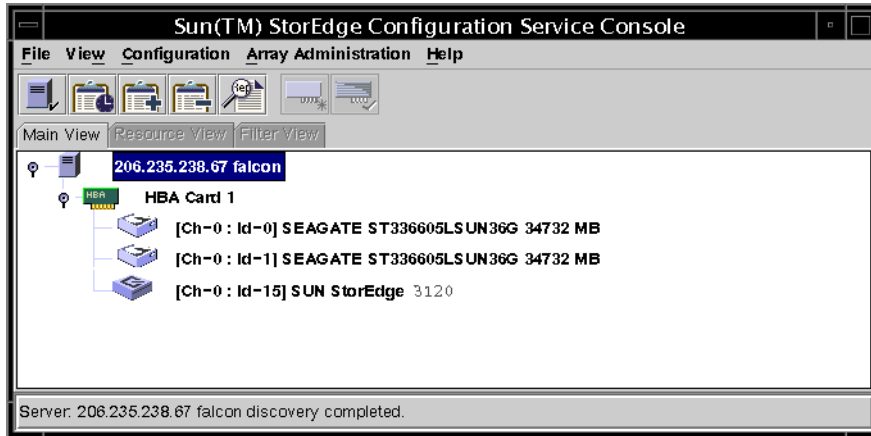
2. Select the Enable JBOD Support check box.

3. To immediately display the JBOD array in the main window, you need to probe for new inventory. Select View → View Server and click Probe.

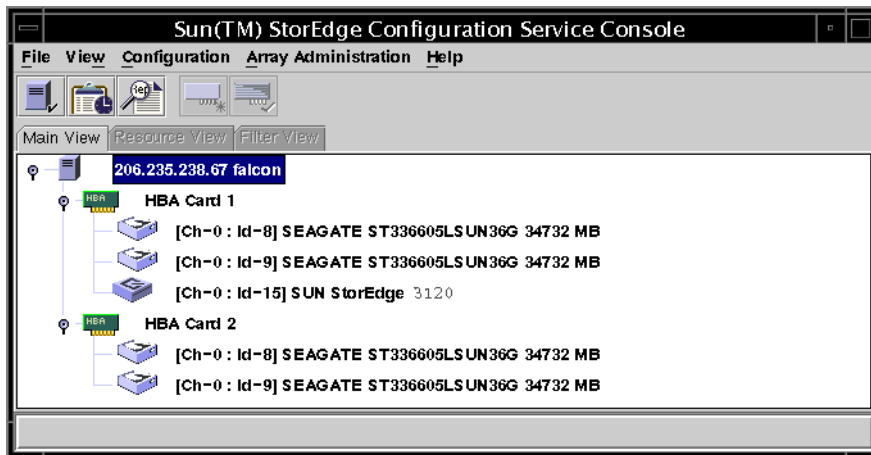
4. Click OK.

The JBOD array is displayed in the main window.

In a single-bus configuration, both ports of the JBOD array are connected to one HBA on the server, as shown in the following example.



In a dual-bus configuration, each port is connected to its own HBA, as shown in the following example.



B.5.2 Sun StorEdge Diagnostic Reporter

Sun StorEdge Diagnostic Reporter supports standalone JBOD arrays. However, triggered event notification is limited to environmental failures and hard drive failures.

B.5.3 Sun StorEdge CLI

The Sun StorEdge CLI supports JBOD arrays. However, because JBOD arrays do not have a RAID controller or RAID controller firmware to manage the disks, this CLI support is limited to the following `sccli` commands:

- `about`
- `download pld-firmware`
- `download ses-firmware`
- `exit`
- `help`
- `quit`
- `select`
- `show frus`
- `show inquiry-data`
- `show pld-revision`
- `show ses-devices`
- `version`

Refer to the “Monitoring JBODs” appendix in the *Sun StorEdge 3000 Family CLI User’s Guide* for information about using these commands.

B.6 Setting Loop IDs on Sun StorEdge 3510 FC JBOD Arrays

When a Sun StorEdge 3510 FC JBOD array is directly attached to a host, unique, hard-assigned loop IDs are assigned to each drive in the JBOD array. A loop ID is the decimal version of an AL_PA. The lowest number loop ID is the lowest priority address on the loop.

On the left front side of a Sun StorEdge 3510 FC JBOD array, underneath the left chassis ear, an ID switch is used to set loop IDs for the disk drives to different ranges of values.

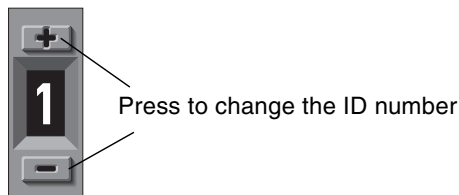


FIGURE B-1 ID Switch

By default, the ID switch on all Sun StorEdge 3510 FC JBOD arrays is set to 0 so the range of IDs is automatically 0 to 11 for 12 drives (the IDs 12-15 are ignored). You can leave the ID switch set to 0, or set it for whichever range of IDs you prefer.

The ID switch offers eight ID ranges. Each set contains 16 IDs (the last 4 IDs in each range are ignored). These ranges are shown in [TABLE B-2](#).

TABLE B-2 ID Switch Settings for Sun StorEdge 3510 FC JBOD Arrays

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

B.6.1 Changing Your ID Switch Setting

Follow these steps if you want to change your ID switch setting:

- 1. Remove the bezel (the front faceplate).**

See [“Replacing the Front Bezel and Ear Caps” on page 7-30](#) for step-by-step instructions for removing and replacing the front bezel and ear caps.



Caution – The plastic ear covers are snap-on parts that require some care when you remove them. Remove the right plastic ear cap carefully to avoid breaking the Reset button directly below the ear cap.

- 2. Remove the left plastic ear cap from the front of the chassis.**
- 3. Use the buttons to change the ID number.**
- 4. Reattach the plastic ear cap and the bezel onto the front of the chassis.**

B.7 Connecting Sun StorEdge 3510 FC JBOD Arrays

You can connect a Sun StorEdge 3510 FC JBOD array directly to a host server. Use your standard host system disk management utilities for all disk management of disks in a Sun StorEdge 3510 FC JBOD array, such as partitioning and formatting. Refer to your host system documentation for more information about disk management.

There are two ways to verify that you have a Sun StorEdge 3510 FC JBOD array (or expansion unit) rather than a Sun StorEdge 3510 FC RAID array:

- The back panel of a Sun StorEdge 3510 FC JBOD array resembles that shown in [“Sun StorEdge 3510 FC JBOD Array Attached to a Single HBA Port”](#) on page B-10.
- Remove the front bezel and look at the labels on the bottom lip of the chassis. If you have a Sun StorEdge 3510 FC JBOD array, one label contains this designation:
3510 AC JBOD (alternating current model) or 3510 DC JBOD (direct current model).

B.7.1 Single-Port Connections to Sun StorEdge 3510 FC JBOD Arrays

You can connect a Sun StorEdge 3510 FC JBOD array to an FC port, a port on a single-port HBA, or a single port on a dual-port HBA.




FIGURE B-2 Sun StorEdge 3510 FC JBOD Array Attached to a Single HBA Port

FIGURE B-2 shows an example of a Sun StorEdge 3510 FC JBOD array with 12 drives connected to a host computer over a single loop to a single FC port. In this example, the port is one of a pair of ports on a dual-port HBA.

In this example, if the loop ID switch is set to 0, the format command shows 12 drives on this loop, ranging from `c1t0d0` to `c1t11d0`.

TABLE B-3 Example of 12 Drives Shown on a 12-Disk Sun StorEdge 3510 FC JBOD Array Connected to a Host Over a Single FC Loop

Drive Identifier			
<code>c1t0d0</code>	<code>c1t3d0</code>	<code>c1t6d0</code>	<code>c1t9d0</code>
<code>c1t1d0</code>	<code>c1t4d0</code>	<code>c1t7d0</code>	<code>c1t10d0</code>
<code>c1t2d0</code>	<code>c1t5d0</code>	<code>c1t8d0</code>	<code>c1t11d0</code>

The drive identifiers are arranged in four columns of three drives each to reflect the physical location of these drives in the array. They correspond to the locations of Disk 1 through Disk 12 as shown in **FIGURE B-5**.

B.7.2 Dual-Port Connections to Sun StorEdge 3510 FC JBOD Arrays

You can also connect a Sun StorEdge 3510 FC JBOD array to both ports on a dual-port HBA, ports on two single-port HBAs, or a port on a single-port HBA and a single port on a dual-port HBA.

Use host-based multipathing software such as VERITAS DMP 3.2 or 3.5 software to manage Sun StorEdge 3510 FC JBOD arrays connected to a host over more than one loop. Multipathing software provides full services on multiple paths between the server and the storage system.




FIGURE B-3 Sun StorEdge 3510 FC JBOD Array Attached to Two HBA Ports

[FIGURE B-3](#) shows an example of a Sun StorEdge 3510 FC JBOD array with 12 drives connected to a host computer over two loops to two FC ports. In this example both ports are on dual-port HBAs.

When you connect your Sun StorEdge 3510 FC JBOD array to a host over two loops, the `format` command shows 24 drives—12 on each loop—because the host sees a device entry for each path. If the loop switch ID is set to 0, these drives are identified as ranging from `c1t0d0` to `c1t11d0` and from `c2t0d0` to `c2t11d0`, as shown in [TABLE B-4](#). However there are, of course, still only 12 physical drives in the JBOD array. These drives are arranged in four columns of three drives each, as shown in [TABLE B-3](#) and [FIGURE B-5](#).

TABLE B-4 Example of 24 Drives Shown on a 12-Disk Sun StorEdge 3510 FC JBOD Array Connected to a Host Over Two FC Loops

Drive Identifier (First Loop)	Drive Identifier (Second Loop)
c1t0d0	c2t0d0
c1t1d0	c2t1d0
c1t2d0	c2t2d0
c1t3d0	c2t3d0
c1t4d0	c2t4d0
c1t5d0	c2t5d0
c1t6d0	c2t6d0
c1t7d0	c2t7d0
c1t8d0	c2t8d0
c1t9d0	c2t9d0
c1t10d0	c2t10d0
c1t11d0	c2t11d0



Caution – Be careful not to mount the same drive twice if you are not using multipathing software.

B.7.3 Changing Your SFP Configuration on JBOD Arrays

Sun StorEdge 3510 FC JBOD arrays use SFP connectors to attach to Fibre Channel ports on your host computer. You might need to move or add one or more SFP connectors to your Sun StorEdge 3510 FC JBOD array to connect it to your host computer.

- To make connections to an empty port, first slide the SFP connector into the port so that it connects firmly with the chassis. Then plug the fiber-optic cable's SFP connector into the duplex jack at the end of the SFP.
- To remove an SFP connector, make sure no cable is connected to it and then slide it out from the port.

For more information about SFPs, refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC Array.

B.7.4 Connecting the Sun StorEdge 3510 FC JBOD Array to a Host Computer

Use fiber-optic cables to connect one or two HBA ports on the host computer to the Sun StorEdge 3510 FC JBOD array, as shown in the previous examples.

1. **Connect a fiber-optic cable to an HBA.**
2. **Set your loop IDs using the ID switch on the JBOD array.**
3. **Connect the SFP connector at the other end of each of these fiber-optic cables to host channel SFP connectors on the back of the JBOD array.**
4. **Power up the equipment in the following order, so that the host computer discovers all connected arrays:**
 - a. **Sun StorEdge 3510 FC JBOD arrays**
 - b. **Host computer**
5. **Use your host system disk management utilities to prepare the disks in the Sun StorEdge 3510 FC JBOD array for use.**

Refer to your host system documentation for more information about disk management utilities available to you.

B.8 Enabling Sun StorEdge 3510 FC JBOD Array Support

For Sun StorEdge 3510 FC JBOD arrays, use JBOD support only when you have a JBOD array connected directly to a host. This enables you to monitor peripheral device condition and events. If you have an expansion unit connected to a RAID array, the RAID controllers monitor the expansion unit's condition and events for you.

Note – Enabling JBOD support can impact I/O performance.

To monitor peripheral device condition and events of a Sun StorEdge 3510 FC JBOD array from the Sun StorEdge Configuration Service console, you first need to enable JBOD support.

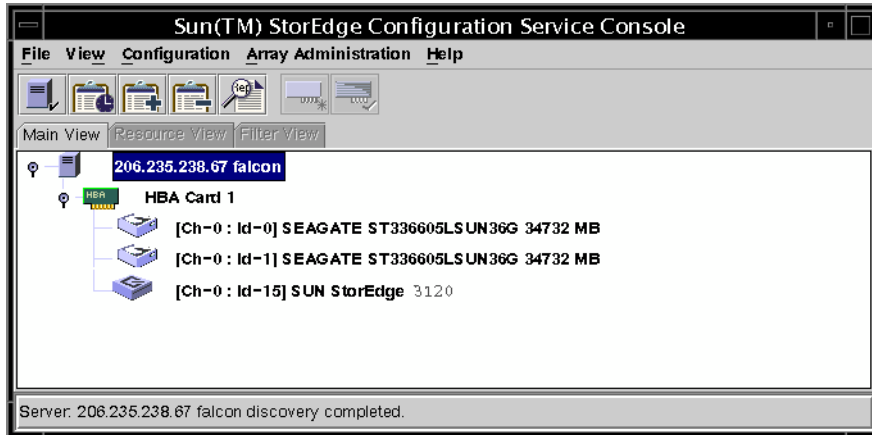
1. **Choose View → Agent Options Management.**

The Agent Options Management window is displayed.

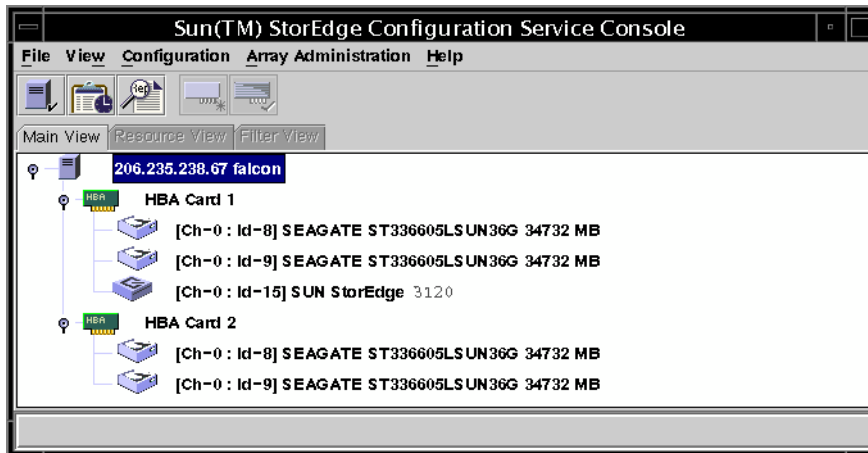
2. Select the Enable JBOD Support checkbox.
3. To immediately display the JBOD array in the main window, you need to probe for new inventory. Choose View → View Server and click Probe.
4. Click OK.

The JBOD array is displayed in the main window.

In a single-bus configuration, both ports of the JBOD array are connected to one HBA on the server, as shown in the following example:



In a dual-bus configuration, each port is connected to its own HBA, as shown in the following example:



B.9 Replacing a Disk Drive in a Sun StorEdge 3510 FC JBOD Array

Replacing a disk drive is a straightforward operation that involves first removing the defective disk drive and then installing a replacement drive.

Note – If you are using disk management software or volume management software to manage your disk storage, you might need to perform software operations to take a disk offline before you remove it, and after you have replaced a drive to bring it back online. Refer to the documentation that accompanies your disk management software or volume management software for more information.

The drive module is hot-swappable; you can replace it while the array is powered on. To replace a disk drive module, first follow the procedure for removing a disk drive and then follow the procedure for installing a new disk drive.

Note – If you are removing a disk drive that you do not intend to replace, install an air management sled in its place. See [“Installing an Air Management Sled” on page B-16](#) for more information.




FIGURE B-4 Front View of a Drive Module Pulled Out of the Chassis

B.9.1 Removing a Disk Drive

To remove a disk drive from an array, perform the following steps.

1. Perform any host software operations necessary to take the disk offline.
2. Unlock the locks with the provided key, and gently pull the plastic front bezel away from the front of the unit so that it drops down and is supported by the two hinged brackets on the sides.
3. Turn the thumbscrew counterclockwise several full turns until the thumbscrew and drive handle are loosened.
4. Gently pull the drive handle upward.
5. Slide the drive module out of the array.

B.9.2 Installing a New Disk Drive

To install a disk drive in an array, perform the following steps.

1. Gently slide the drive module into the drive slot until the handle pins slip into the chassis notch.
2. Lower the disk drive handle until it is vertical.
3. Press and hold the drive handle in while you press the thumbscrew in until it engages the threads.
4. Turn the thumbscrew clockwise until it is finger-tight.

Note – To ensure that a thumbscrew is finger-tight, tighten it with a screwdriver and then loosen the thumbscrew counterclockwise a quarter-turn.

5. Push the plastic front bezel onto the front of the unit until it is seated firmly, and use the key to lock the locks.
6. Perform any operations your host software requires to recognize the new drive and bring it under software control.

B.9.3 Installing an Air Management Sled

The front of an air management sled looks similar to a disk drive module, but the air management sled is an empty box used to maintain optimum airflow in a chassis.

If you have removed a disk drive and do not intend to replace it, insert an air management sled to maintain the optimum airflow inside the chassis. You can install the air management sled by using the same procedure as [“Installing a New Disk Drive” on page B-16](#), except that no software operations are necessary.

B.10 Downloading Firmware to Disk Drives in a JBOD

For instructions on how to download firmware to disk drives in a JBOD directly attached to a host, refer to the README file in the patch that contains the firmware.

B.11 Understanding Component Status LEDs on JBOD Arrays

This section describes the meaning of front panel and back panel LEDs you can use to determine the status of the operation of all drives and modules.

B.11.1 LED Activity When JBOD Arrays Are Powered On

It can be useful to power on an array before it is connected to a host. If a Sun StorEdge 3510 FC JBOD array is powered on and not connected to a server, the LED activity shown in the following table is visible.

TABLE B-5 Front Panel LED Status When Array Is First Powered On

Drive LEDs	Solid green
Chassis ear LEDs	Solid green

B.11.2 Front Panel LEDs on JBOD Arrays

The drive LEDs are located between the rows of drives on the front panel, as shown in [FIGURE B-5](#).




FIGURE B-5 Front Panel LEDs

The system operation LEDs and Reset button are located on the right ear of the chassis, as shown in [FIGURE B-6](#). Use a paperclip to push the Reset button to silence a failed component alarm. See [“Silencing Audible Alarms” on page 7-5](#) for more information about silencing audible alarms.




FIGURE B-6 Chassis Ear LEDs and Reset Button on Front Panel

TABLE B-6 shows the JBOD array's front panel LEDs and their definitions.

TABLE B-6 JBOD Front Panel LEDs

LED	LED Color	Description
Drive	Solid green	Good: Drive power-up and spin-up OK.
	Blinking green	Good: Drive activity.
	Solid amber	Failed: Drive failure.
Power (Light bulb icon) Monitors the DC output voltage within tolerance specification. Overcurrent protection shutting down any voltage output is also displayed. Voltage thresholds: +5 VDC +/- .25 VDC +12 VDC +/- .6 VDC Current thresholds: +5 VDC 35A +12 VDC 25A	Solid green	Good: Power supply good.
	Solid amber	Failed: One or more output voltages out-of-range.
Fan (Fan icon) Monitors the fan speed within nominal operational RPM specification of 5000 RPM.	Solid green	Good: Both fans rotating at 4000 RPM or higher.
	Solid amber	Faulty/Failed: Faulty/Failed: One or both fans rotating a less than 4000 RPM.
Temp (Thermometer icon) Monitors the temperature level and indicates violations of internal temperature thresholds of 55C and 60C.	Solid green	Good: Under temperature threshold of 55C.
	Solid amber	Failed: Temperature threshold equal to or over 55C.
	Blinking amber	Failed: Temperature threshold equal to or over 60C. Blinking frequency of 4 Hz +/- 1 Hz.
Event (Caution icon) Indicates any abnormal or failure event in the I/O board.	Solid green	Normal operation of the I/O board.
	Solid amber	Failed I/O board.
	Blinking amber	Indicates that the version of the SES firmware or associated hardware PLD code on one controller does not match that on the other controller.

B.11.3 Back Panel LEDs on JBOD Arrays

The back panel LED colors indicate the conditions described in the following figures and tables.

Note – Although an amber light on the back panel often indicates a faulty component, a steady amber light on the Ethernet Link LED indicates normal Ethernet activity. See [TABLE B-7](#) for more information.

B.11.4 Back Panel I/O Module LEDs

[FIGURE B-7](#) shows the I/O module and its LEDs.




FIGURE B-7 I/O Module for a Sun StorEdge 3510 FC JBOD Array

The following table shows the I/O module LEDs and their color definitions.

TABLE B-7 JBOD I/O Controller Module and Battery Module LEDs

LED	Purpose	LED Color Definition
Activity	I/O activity of host and disk ports	Off – Not busy (no I/O activity) Blinking green – Busy (active I/O)
SFP Link (L)	SFP link status	Solid green – Active good FC Connection Off – Empty or failed FC connection
SFP Speed (S)	SFP speed status	Solid green – 2 Gbit Off – 1 Gbit

Note – An SFP Link Status (L) LED is off if there is no connection, or a failed connection, to this SFP.

B.11.5 Back Panel Power Supply and Fan Module LEDs on JBOD Arrays

The following figure shows the AC power supply and fan module.




FIGURE B-8 AC Power Supply and Fan Module

The following figure shows the DC power supply and fan module.




FIGURE B-9 DC Power Supply and Fan Module

TABLE B-8 shows the color definitions for the power supply LED.

TABLE B-8 Power Supply LEDs

Purpose	LED	LED Color Definition
Monitors the DC output voltage within tolerance specification. Overcurrent protection shutting down any voltage output is also displayed.	Solid green	Power supply and fans are good.
Voltage thresholds: +5 VDC +/- .25 VDC +12 VDC +/- .6 VDC Current thresholds: +5 VDC 35A +12 VDC 25A	Solid amber	Failed: One or more output voltages out-of-range or fan speed is under 3150 RPM.

B.12 Troubleshooting Sun StorEdge 3510 FC JBOD Arrays

This section describes troubleshooting procedures and error messages you can use to isolate configuration and hardware problems. For additional troubleshooting information refer to the “Troubleshooting Your Array” chapter in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for the Sun StorEdge 3510 FC Array.

B.12.1 Troubleshooting Configuration Issues

Follow this sequence of general steps to isolate software and configuration issues.

1. **Look for storage-related messages in `/var/adm/messages` and identify any suspect Sun StorEdge 3510 FC JBOD arrays and fibre loops.**
2. **Check your Sun StorEdge Configuration Service console for alerts or messages.**
3. **Check revisions of software package, patches, and hardware.**
4. **Verify correct device file paths.**
5. **Check any related software, configuration, or startup files for recent changes.**

Search SunSolve Online for any known related bugs and problems at:
<http://sunsolve.Sun.COM>

B.12.2 Troubleshooting Hardware Issues

When a problem is not otherwise reproducible, suspect hardware might need to be replaced. Always make only one change at a time and carefully monitor results. When possible, it is best to restore the original hardware before replacing another part to eliminate the introduction of additional unknown problem sources.

After hardware replacement, a problem can usually be considered solved if it does not resurface during a period equal to twice its original frequency of occurrence. For example, if a problem was occurring once a week on average before a potential fix was made, running two weeks without seeing the problem again suggests a successful fix took place.

Troubleshooting hardware problems is usually accomplished by an FRU isolation sequence that uses the process of elimination. Set up a minimal configuration that shows the problem and then replace elements in this order, testing after each replacement until the problem is solved:

- Replace the cables.
- Replace the SFPs.
- Replace the drives.
- Replace the HBAs.

Often you can also find out what does cause a hardware problem by determining the elements that do not cause it. Start out by testing the smallest configuration that does work, and then keep adding components until a failure is detected.

To view error messages reported by JBODs or expansion units, use any of the following:

- Message log in the operating system
 - For Solaris, `var/adm/messages`
 - For Linux, `var/log/message`
 - For Microsoft Windows, use the `event viewer`
- Sun StorEdge Configuration Service event log
- Sun StorEdge CLI event log
- RAID firmware event log for expansion units

For more information about replacing the chassis, I/O module, or controller, refer to the *Sun StorEdge 3000 Family FRU Installation Guide for 2U Arrays*.



Caution – When you rotate or replace drives, make sure that:

- All I/O is stopped.
- The “Periodic Drive Check Time” setting in the firmware application is set to disabled (this is the default setting). This prevents automatic drive rebuild, which is not recommended for live systems or troubleshooting.

For more information, refer to “Periodic Drive Check Time” in the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for your array.



Caution – To prevent any possible data loss, back up the chassis data onto another storage device prior to replacing a disk drive.

Before you begin troubleshooting JBODs or expansion units, check the cables that connect the host to the JBOD or expansion unit. Look for bent pins, loose wires, loose cable shields, loose cable casing, and any FC cables with 90 degree or more bends in them. If you find any of these problems, replace the cable.

The [FIGURE B-10](#) flowchart provides troubleshooting procedures specifically for JBODs and expansion units. For additional troubleshooting flowcharts, see “[Power Supply and Fan Module](#)” on page 8-16 and “[Drive LEDs](#)” on page 8-19.

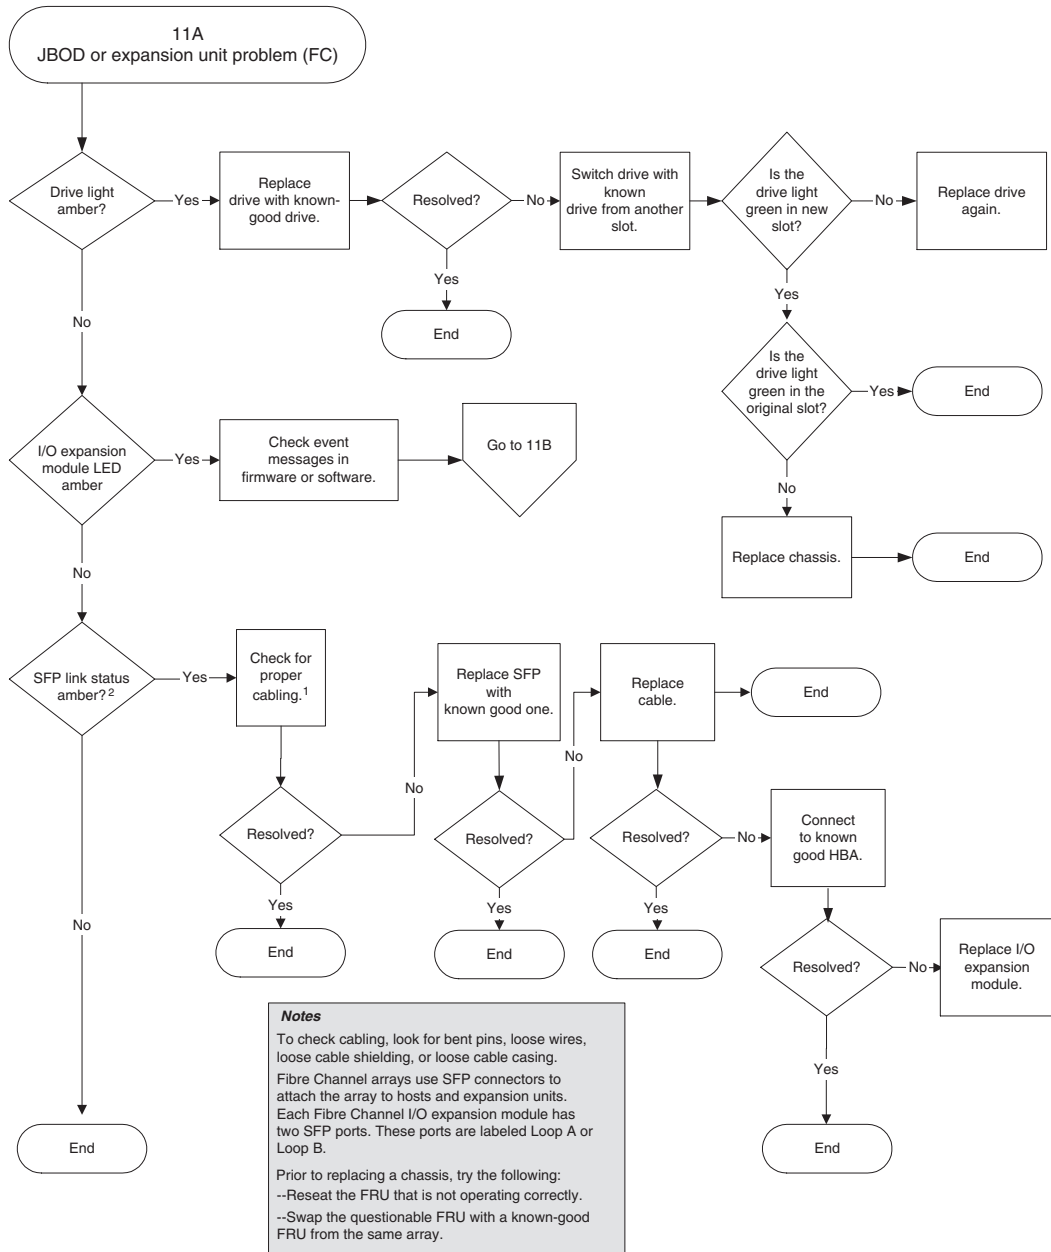


FIGURE B-10 JBOD or Expansion Unit Flowchart 1 of 2

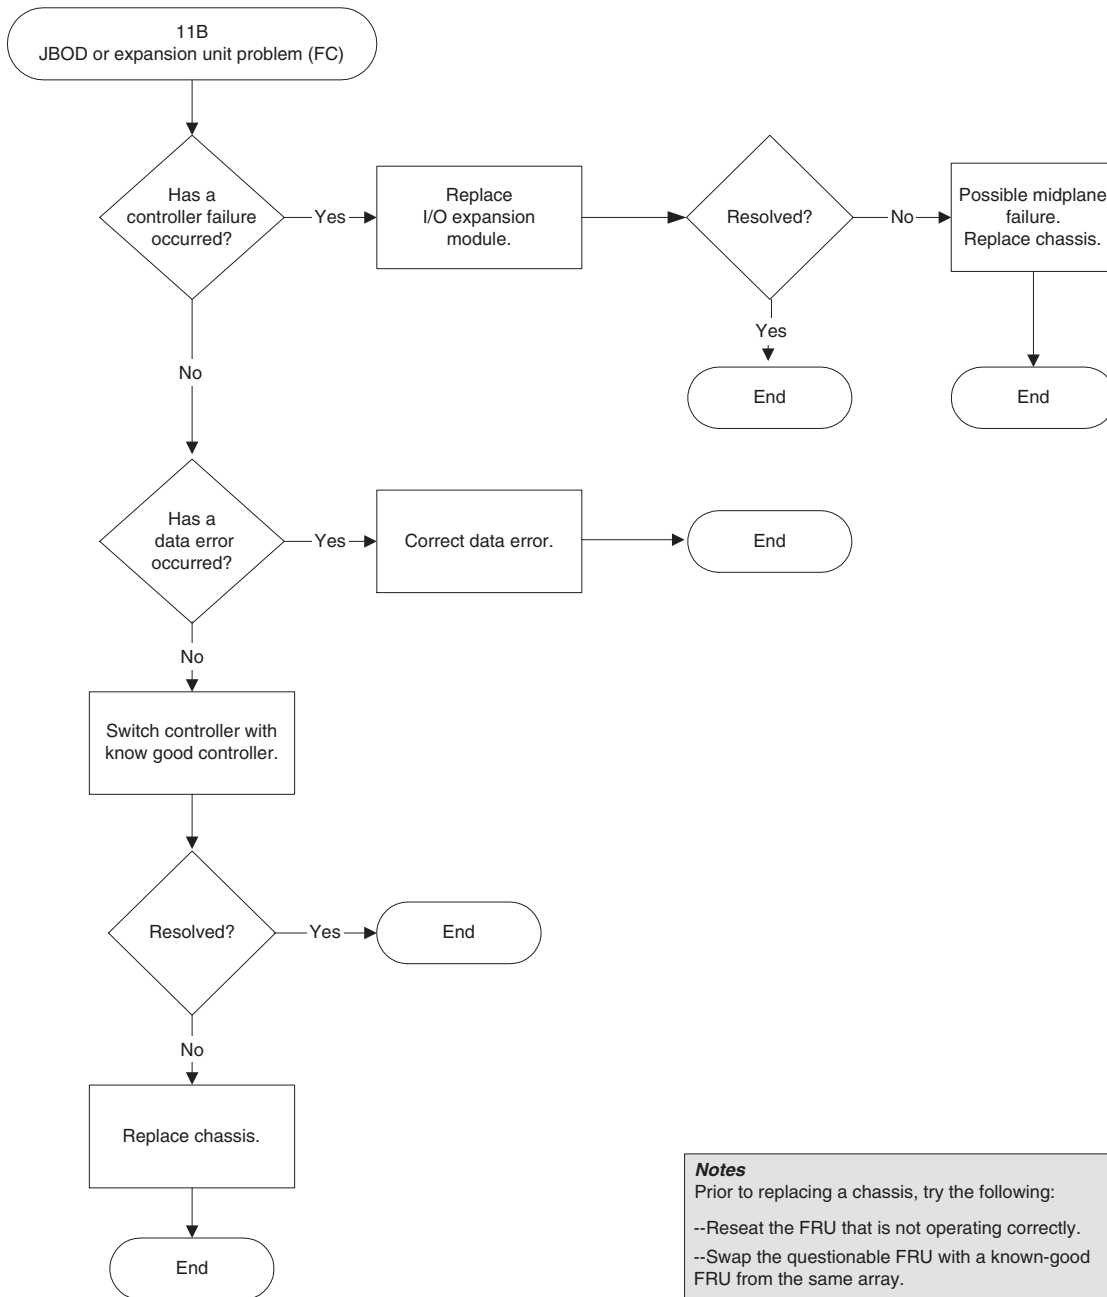


FIGURE B-11 JBOD or Expansion Unit Flowchart 2 of 2

B.13 Converting JBODs to RAID arrays

If you have an FC JBOD and want to convert it to an FC RAID array, refer to the *Sun StorEdge 3000 Family FRU Installation Guide* for instructions.

Failed Component Alarm Codes

This appendix describes failed component alarm codes.

Failed component alarm codes use Morse code dot and dash characters. The dot “.” is a short tone sounding for one unit of time. The dash “-” is a long tone sounding for three units of time.

Alarms, also referred to as beep codes, are presented in a sequence, starting with the critical component failure alarm, which alerts you to a component problem or failure or a firmware mismatch. This alarm is then followed by alarms for whichever components or assemblies have failed. Once the beep code sequence is complete, it repeats. To understand the beep codes, listen to the sequence of codes until you can break down the sequence into its separate alarms. You can also check your software or firmware for alarms, error messages, or logs to isolate and understand the cause. For example, in the case of a fan failure in a power supply, you might first hear the critical component failure alarm, followed by a power supply failure alarm from power supply 0 or power supply 1, followed by a fan failure event alarm, followed by an event alarm. This sequence continues to repeat.

TABLE C-1 Failed Component Alarm Codes

Failure	Morse Code Letter	Morse Code Sound Pattern
Critical component failure or mismatch	8 dashes	-----
Power supply 0 failure	P0	. -- . -----
Power supply 1 failure	P1	. -- . . -----
Event alarm	E	.
Fan failure	F	. . - .

TABLE C-1 Failed Component Alarm Codes (*Continued*)

Failure	Morse Code Letter	Morse Code Sound Pattern
Voltage failure	V	. . . -
Temperature failure	T	-
SES/PLD firmware mismatch	R	. - .

See “[Front Panel LEDs](#)” on page 6-1 for more information about the SES/PLD firmware mismatch mentioned in [TABLE C-1](#).



Caution – Be particularly careful to observe and rectify a temperature failure alarm. If you detect this alarm, shut down the controller and the server as well if it is actively performing I/O operations to the affected array. Otherwise, system damage and data loss can occur.

Record of Settings

This appendix provides tables for recording configuration data. Topics covered are:

- [“View and Edit Logical Drives” on page D-2](#)
- [“View and Edit Host LUNs” on page D-4](#)
- [“LUN Mappings” on page D-4](#)
- [“View and Edit SCSI Drives” on page D-5](#)
- [“View and Edit SCSI Channels” on page D-6](#)
- [“View and Edit Peripheral Devices” on page D-7](#)
- [“Save NVRAM to Disk and Restore From Disk” on page D-7](#)

In addition to saving the configuration data in NVRAM to disk, keeping a hard copy of the controller configuration is also recommended. This speeds the re-creation of the RAID in the event of a disaster.

The following tables are provided as a model for recording configuration data.

Note – The configuration data in the NVRAM should be saved to disk or file whenever a configuration change is made.

D.1 View and Edit Logical Drives

D.1.1 Logical Drive Information

LG	ID	LV	RAID Level	Size (MB)	Status	O	#LN	#SB	#FL	Name	Disk Reserved Space

D.3 View and Edit SCSI Drives

Slot	Channel	ID	Size (MB)	Speed	LG DRV?		Vendor and Product ID
					Global Spare?	Local Spare?	

Cable Pinouts

This appendix identifies the pinouts for each connector used in the Sun StorEdge 3510 FC Array. Topics covered are:

- “RJ-45 Connector” on page E-1
- “DB-9 COM port” on page E-2

E.1 RJ-45 Connector




FIGURE E-1 Ethernet RJ-45 Socket 10/100 BASE-T

TABLE E-1 Ethernet RJ-45 Pin Description

Pin No.	Description.	Color
1	TX +	White with orange
2	TX -	Orange
3	RX +	White with green
4		Blue
5		White with blue

TABLE E-1 Ethernet RJ-45 Pin Description (*Continued*)

Pin No.	Description.	Color
6	RX –	Green
7		White with brown
8		Brown

E.2 DB-9 COM port

The COM port is a female DB-9 connector that requires a male DB-9 null modem cable.




FIGURE E-2 RS-232 DB9 (EIA/TIA 574) View of the Male End

TABLE E-2 Pin Names

Pin No.	Name	Notes/Description
1	DCD	Data Carrier Detect
2	RD	Receive Data (also called RxD, Rx)
3	TD	Transmit Data (also called TxD, Tx)
4	DTR	Data Terminal Ready
5	SGND	Ground
6	DSR	Data Set Ready
7	RTS	Request To Send
8	CTS	Clear To Send
9	RI	Ring Indicator

Configuring a Sun Server Running the Solaris Operating Environment

This appendix provides access information and LUN setup information needed when you connect a Sun StorEdge 3510 FC Array to a Sun server running the Solaris operating environment. For a list of supported adapters, refer to the *Sun StorEdge 3510 FC Array Release Notes*.

This information supplements the configuration steps presented in this manual and covers the following topics:

- [“Setting Up the Serial Port Connection” on page F-1](#)
- [“Accessing the Firmware Application From a Solaris Host” on page F-2](#)
- [“Using the `tip` Command for Local Access to the Array” on page F-3](#)
- [“Determining the WWN in the Solaris Operating Environment” on page F-3](#)

F.1 Setting Up the Serial Port Connection

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

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program once you have assigned an IP address to the array. For details, see [“Configuring a COM Port to Connect to a RAID Array” on page 4-14](#) and refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

1. Use a null modem serial cable to connect the COM port of the RAID array to an unused serial port on a Solaris workstation.

A DB-9-to-DB-25 serial cable adapter is included in your package contents to connect the serial cable to a DB-25 serial port on a workstation.




FIGURE F-1 RAID Array COM Port Connected Locally to the Serial Port of a Workstation

2. Power up the array.

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

3. Set the serial port parameters on the workstation.

See [“Configuring a COM Port to Connect to a RAID Array”](#) on page 4-14 for the parameters to use.

F.2 Accessing the Firmware Application From a Solaris Host

See [“Using the `tip` Command for Local Access to the Array”](#) on page F-3 for setting serial port parameters on a Solaris host.

See [“Using the `tip` Command for Local Access to the Array”](#) on page F-3 for instructions on using the Solaris operating environment’s `tip` command.

F.3 Using the `tip` Command for Local Access to the Array

You can access the RAID COM port (DB-9 connector on the controller module) locally with the following steps.

1. Connect the RAID array COM port to the serial port of a Solaris workstation, as shown in [FIGURE F-1](#). Use the `tip` command to access the array locally:

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

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

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

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

F.4 Determining the WWN in the Solaris Operating Environment

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

```
# luxadm probe
```

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

A terminal window titled "Terminal" with a menu bar containing "Window", "Edit", "Options", and "Help". The terminal shows the command "luxadm probe" being executed on a "falcon" system. The output lists two Fibre Channel devices with their Node WWNs, Device Types, and Logical Paths.

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

FIGURE F-2 Worldwide Name Information Displayed by the luxadm command.

Configuring a Windows 2000 Server or Windows 2000 Advanced Server

This appendix provides platform-specific host installation and configuration information to use when you connect a Sun StorEdge 3510 FC array to a host running the Windows 2000 Server or Windows 2000 Advanced Server operating system.

The Sun StorEdge 3510 FC Array supports Microsoft Windows 2000 Server and Microsoft Windows 2000 Advanced Server, Service Pack 3 or greater, in dual-path configurations using Sun StorEdge Traffic Manager 3.0 failover drivers for Microsoft Windows 2000.

If you are going to use dual-path instead of single-path configuration, refer to the *Sun StorEdge Traffic Manager 3.0 Installation and User's Guide for Microsoft Windows 2000 and Windows NT* for detailed instructions about setting up the device driver on the server and for additional information about configuring your Windows 2000 Server or Advanced Server.

Customers interested in Sun StorEdge Traffic Manager 3.0 for multiplatform support should contact Sun Sales or visit:

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

For more information on multiplatform support, see:

http://www.sun.com/storage/san/multiplatform_support.html

The information in this appendix supplements the configuration steps presented in [Chapter 5](#), and covers the following steps:

- [“Setting Up the Serial Port Connection” on page G-2](#)
- [“Accessing the Firmware Application From a Windows 2000 Server or Windows 2000 Advanced Server” on page G-5](#)
- [“Enabling a Windows 2000 Server or Windows 2000 Advanced Server to Recognize New Devices and LUNs” on page G-5](#)

- [“Determining the World Wide Name for Windows 2000 Servers and Windows 2000 Advanced Servers”](#) on page G-10

G.1 Setting Up the Serial Port Connection

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

Note – You can also monitor and configure a RAID array over an IP network using the RAID firmware or Sun StorEdge Configuration Service software after you assign an IP address to the array. For details, refer to [“Setting Up Out-of-Band Management Over Ethernet”](#) on page 4-16 and to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

1. **Use a null modem serial cable to connect the COM port of the RAID array to an unused serial port on your host system.**

A null modem cable has serial signals swapped for connecting to a standard serial interface.

Note – A DB-9-to-DB-25 serial cable adapter is included in your package contents to connect the serial cable to a DB-25 serial port on your host if you do not have a DB-9 serial port.

2. **Power up the array.**
3. **On the server, choose Start → Programs → Accessories → Communications → HyperTerminal.**
4. **Type a name and choose an icon for the connection.**

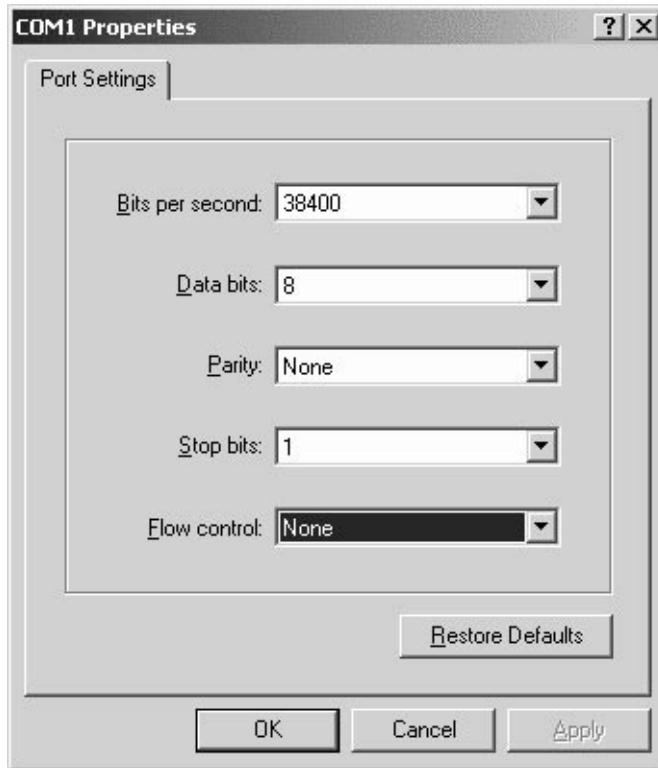
5. In the Connect To window, select the COM port from the “Connect using:” drop-down menu for that array.



6. Click OK.

7. In the Properties window, set the serial port parameters using the drop-down menus.

See “[Configuring a COM Port to Connect to a RAID Array](#)” on page 4-14 for the parameters to use.



8. To save the connection and its settings, choose File → Save.

The connection file name is *connection-name* where *connection-name* is the name you gave this HyperTerminal connection when you created it.

9. To make a connection shortcut on your desktop, choose Start → Find → For Files or Folders. Enter *connection-name* and click the Search Now button. Highlight and right-click on the file name in the Search Results window, choose Create Shortcut, and click Yes.

You are now ready to configure your array.

G.2 Accessing the Firmware Application From a Windows 2000 Server or Windows 2000 Advanced Server

To access the array from a Windows 2000 Server or Windows 2000 Advanced Server through a serial port, use the HyperTerminal connection you configured in [“Setting Up the Serial Port Connection”](#) on page G-2 or install a VT100 terminal emulation program and use the parameters described there.

To access the array through an Ethernet connection, follow the procedures in [“Setting Up Out-of-Band Management Over Ethernet”](#) on page 4-16. To access the command prompt described there from a Windows 2000 Server or Windows 2000 Advanced Server, choose Programs → Accessories → Command Prompt.

G.3 Enabling a Windows 2000 Server or Windows 2000 Advanced Server to Recognize New Devices and LUNs

Before beginning this procedure, make sure that you are using a supported FC host bus adapter (HBA) such as a QLogic QLA2310 or QLogic QLA2342. Refer to the release notes for your array for current information about which HBAs are supported.

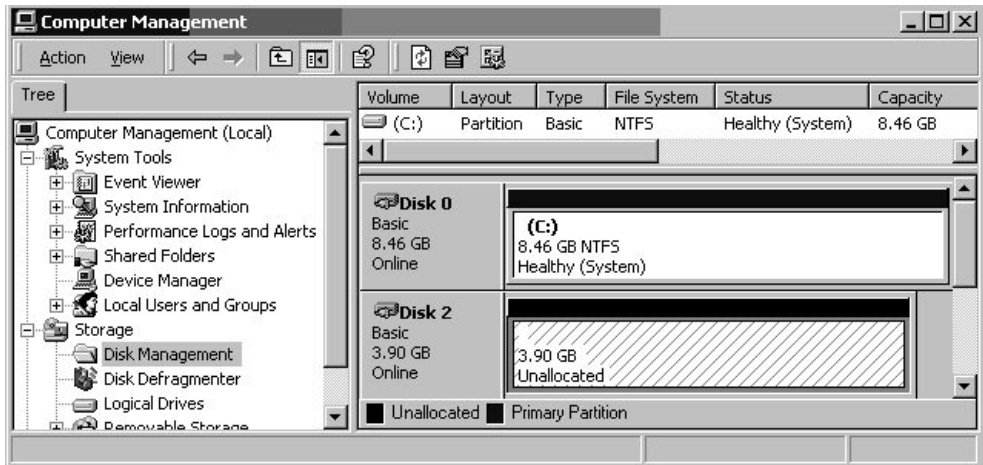
Also make sure that you are using a supported driver for your HBA. For the QLA2310 or QLA2342, use driver version 8.1.5.13 or later.

Note – When the Microsoft Windows operating system recognizes the Sun StorEdge 3510 FC Array’s SCSI Enclosure Services (SES) processor and identifies it as a SE3510 Scsi Enclosure, it might prompt you to supply a driver. This is standard plug-and-play behavior when unknown devices are encountered, but no device driver is necessary. If you see this prompt, simply press Cancel. If you have multiple arrays, you might see this prompt more than once. Each time you see this prompt, press Cancel.

1. **Boot your system and verify that the HBA basic input/output system (BIOS) recognizes your new FC device.**

Note – While your system is starting up, you should see your new Fibre Channel devices.

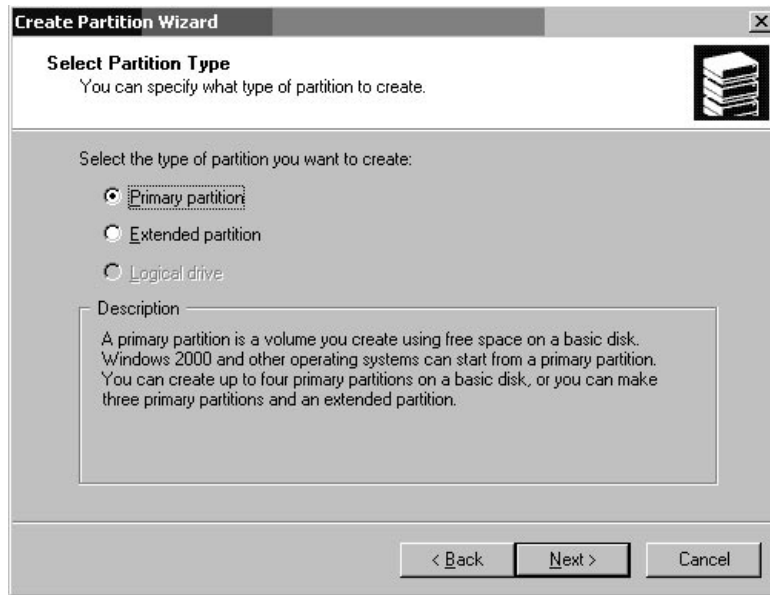
2. If a Found New Hardware Wizard is displayed, click **Cancel**.
You are now ready to partition and format your new devices.
3. Open the Disk Management folder.
 - a. Right-click on the My Computer icon to display a pop-up menu.
 - b. Choose **Manage**.
 - c. Select the Disk Management folder.
 - d. If a Write Signature and Upgrade Disk Wizard is displayed, click **Cancel**.
A “Connecting to Logical Disk Manager Server” status message is displayed.
4. Select your new device when it is displayed.



5. Right-click in the Unallocated partition of your device to display a pop-up menu.

6. Choose Create Partition from the pop-up menu.


The Create Partition wizard is displayed.



7. Click Next.

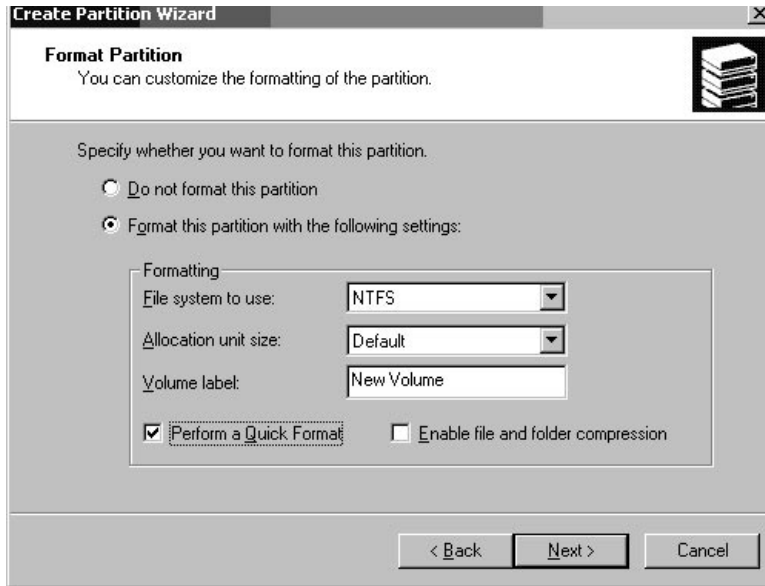
8. Select Primary partition and click Next.

9. Specify the amount of disk space to use or accept the default value, and click Next.



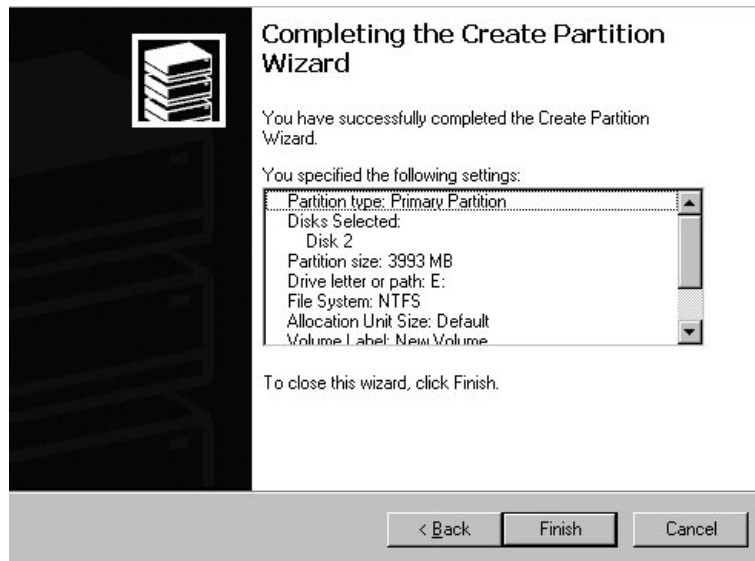
10. Assign a drive letter, if desired, and click Next.
11. Select "Format this partition with the following settings."
 - a. Specify NTFS as the "File system to use."

b. Make sure the “Perform a Quick Format” checkbox is selected.




c. Click Next.

A confirmation dialog box displays the settings you have specified.



12. Click Finish.

The new partition is formatted and the formatted partition is identified as NTFS in the Computer Management window.



13. Repeat these steps for any other new partitions and devices you want to format.

G.4 Determining the World Wide Name for Windows 2000 Servers and Windows 2000 Advanced Servers

Before you can create host filters, you need to know the world wide name (WWN) for the FC HBA that connects your host to your FC array.

- 1. Boot a specific host system and note the BIOS version and HBA card models connected to your host.**
- 2. Access the HBA card's BIOS with the appropriate command (alt-q or control-a are commonly used).**

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

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

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

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

See [“Setting Up Host Filter Entries”](#) on page 5-52 for more information about creating host filters.

Configuring a Linux Server

This appendix provides access information and LUN setup information needed when you connect a Sun StorEdge 3510 FC Array to a server running the Linux operating environment. For a list of supported adapters, refer to the *Sun StorEdge 3510 FC Array Release Notes*.

This information supplements the configuration steps presented in this manual and covers the following topics:

- [“Setting Up the Serial Port Connection” on page H-1](#)
- [“Checking the Adapter BIOS” on page H-4](#)
- [“Multiple LUN Linux Configuration” on page H-5](#)
- [“Making an ext3 File System for Linux” on page H-6](#)
- [“Creating a File System” on page H-7](#)
- [“Creating a Mount Point and Mounting the File System Manually” on page H-7](#)
- [“Mounting the File System Automatically” on page H-8](#)
- [“Determining the World Wide Name for Linux Hosts” on page H-9](#)

H.1 Setting Up the Serial Port Connection

The RAID controller can be configured by means of a host system running a VT100 terminal emulation program or running a terminal emulation program such as Minicom.

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program after you assign an IP address to the array. For details, see [“Setting Up Out-of-Band Management Over Ethernet” on page 4-16](#) and refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

To access the controller firmware through the serial port, perform the following steps:

1. **You must use a null modem cable to communicate via the serial port of the array. Connect the serial null modem cable to the array and to ttyS0 (Com1), ttyS1 (Com2), ttyS2 (Com3), or ttyS3 (Com4).**

Note – A DB-9-to-DB-25 serial cable adapter is included in your package contents to connect the serial cable to a DB-25 serial port on your host if you do not have a DB-9 serial port.

2. **Power on the array.**
3. **After the array is powered up, power on the Linux server, and log in as root (or log in as su if you logged in as a user).**
4. **Open a terminal session and type:**

```
minicom -s
```

Press Return. You see the setup menu, where you define which serial ports to use, baud rate, hand shake settings, and flow control.

5. **Set the serial port parameters on the server.**

See [“Configuring a COM Port to Connect to a RAID Array”](#) on page 4-14 for the parameters to use.

- a. **At the configuration screen, highlight Serial Port Settings using the arrow keys and press Return.**
- b. **If A is not correct, press the letter A and the cursor goes to line A.**

Backspace and edit:

```
/dev/ttyS(x)
```

where *x* is the proper serial port connecting the server to the array.

After you edit the line, press Return and the cursor goes back to the line Change which Settings?

- c. If E is not correct, press the letter E and the cursor goes to line E. Backspace and change the line to:

```
38400 8N1
```

After you edit the line, press Return and the cursor goes back to the line Change which Settings?

- d. Set F to no. When pressing F, toggle from yes to no or vice versa.
- e. Set G to no. When pressing G, toggle from yes to no or vice versa.
- f. Press the Esc key to return to the configuration screen.
- g. In the configuration menu, highlight "Save setup as df1" using the arrow keys and press Return.
A "configuration saved" confirmation message is displayed.
- h. Highlight "Exit from Minicom" in the configuration menu, and press Return.

H.2 Accessing the Firmware Application From a Linux Server

After you have set the Minicom serial port parameters, use the Minicom utility to access the controller firmware with the following steps:

1. To start the Minicom program from the terminal window, type:

```
minicom
```

A pop-up window stating "Initializing Modem" appears. When initialization is complete, the following banner is displayed:

```
Welcome to Minicom 2.00.0
```

```
Options: History Buffer, F-Key Macros, Search History Buffer, I18n
```

```
Compiled on Jun 23 2002, 16:14:20.
```

```
Press "CTRL-A" Z for help on special keys.
```

2. Press Ctrl-l to connect to the array and access the firmware application.

3. In the main screen, choose **Terminal (VT100 Mode)** and view the main menu commands used to configure the array.

H.3 Checking the Adapter BIOS

When booting the server, watch for the HBA card BIOS message line to show up. Then press the proper sequence of keys to get into the HBA BIOS:

- For FC Qlogic HBAs, <Alt-Q>

The keystrokes are listed on the screen when the adapter is initializing.

After you enter the Qlogic HBA BIOS with Alt-Q, perform the following steps.

If your HBA has more than one interface on it, they show up in the main screen of the BIOS software after you press Alt-Q. You have to make the same changes on one interface as the other unless one of the interfaces is going to be bootable.

1. **If you have more than one interface, highlight the top one listed under Adapter Type and press Return. If you only have one interface, it is already highlighted; press Return.**
2. **Highlight Configuration Setting and press Return.**
3. **Highlight Host Adapter Settings and press Return.**
4. **Move down to Host Adapter BIOS and make sure it is enabled; if not, press the Return key and it toggles from disabled to enabled. The setting should be enabled.**
5. **Press Esc and go back to Configuration Settings. Highlight Selectable Boot Settings and press Return.**

This is where you can make the interface bootable or not bootable.

6. **Highlight Select Boot Device. Press the Return key to toggle from disabled to enabled.**
 - Select disabled if this is not going to be a bootable device.
 - If it is going to be bootable device, select Enabled.
7. **Press Esc until you get back to Configuration Settings.**
8. **Highlight Extended Firmware Settings and press Return.**

9. **On the Extended Firmware Settings menu, highlight Connection Option and press Return. A screen lists three types of connection:**
 - 0 – Loop only- - - for connecting to loop type devices
 - 1 – Point to point only - - - connecting to fabric switches
 - 2 – Loop preferred, otherwise point to point
10. **If you are connecting to a loop device, highlight 0 and press Return. If you are connecting to a Fabric device or switch, highlight 1 and press Return.**
11. **Once completed, press Esc until you get to a screen that displays “Configuration setting modified.” Highlight Save changes and press Return.**
12. **Return to the Fast!UTIL Options. Highlight Scan Fibre Devices and press Return.**

This menu option scans all 126 channels to see if there are any devices attached; the devices show up after the scan. If there are no devices attached, it takes some time to scan. If there are devices attached, it usually finds them right away.
13. **If you are satisfied with the configuration, press Esc until you get to Configuration Settings.**
14. **Highlight Exit Fast!UTIL and press Return. A screen displays Exit Fast!UTIL. Highlight Reboot System and press Return.**

The server reboots.

H.4 Multiple LUN Linux Configuration

By default, the Linux kernel does not support multiple LUNs. To support multiple LUNs, modify the kernel with the following steps:

1. **Log in as root, or su to root if you are logged in as a user.**
2. **Add this line to the end of the `/etc/modules.conf` file and save the file:**

```
options scsi_mod max_scsi_luns=255
```

3. **At the system prompt, enter this command and press Return:**

```
mkinitrd -f /boot/initrd-2.4.9-e.3.img 2.4.9-e.3
```

The 2.4.9-e.3 entries refer to the current kernel. To find out your current kernel, type `uname -r` and substitute your kernel information in place of the 2.4.9-e.3 entries.

4. Reboot the server.

To halt the server completely, use `shutdown -h now`. To reboot automatically after the shutdown is complete, use `shutdown -r now`.

H.5 Making an ext3 File System for Linux

The following procedure to label and partition drives using `fdisk` applies to an ext3 file system. To discover which disk you want to label, you need to find out what device it is.

1. To list all devices and their paths, start a terminal session and type:

```
fdisk -l
```

Record the device names and paths that you plan to use.

2. Then type:

```
fdisk /dev/sd(x) x= a,b,c,...
```

A banner appears for the specified array device. The last statement displays a prompt.

3. Type `m` or `help`.

4. On the displayed menu, select “n” for the command action and press Return.

Two choices are displayed:

e extended

p primary partition (1-4)

Only four primary partitions are allowed per array. All additional partitions must be added as extended LUNs under one primary partition. Only one primary partition is allowed to have extended LUNs.

5. For the first partition, select “p.”

When several options appear, keep the defaults. You can reconfigure this after you understand the process and see what it looks like. Add additional primary partitions and extended partitions as needed.

6. After you have completed accepting the defaults and are back at the “Command (m or help):” screen, press **W** to save the configuration and exit `fdisk`.

Your partition is ready for a file system now.

H.6 Creating a File System

1. Log in as root, or `su` to root if you are logged in as a user.
2. Take the device that you ran `fdisk` on and run the following command to create an ext3 file system:

```
mkfs -t ext3 /dev/sdb(x)
```

where `x` is the partition that you are making into a file system. Replace `x` with `1` since there is only one partition.

H.7 Creating a Mount Point and Mounting the File System Manually

1. Go to the directory where you want to create a directory that will be the mount point. Then type the following command:

```
mkdir (name)
```

2. To mount your file system, type the following:

```
mount /dev/sdb(x) / directory-path
```

where `x` is `1` for this partition and the *directory-path* is the directory that was created and its location.

H.8 Mounting the File System Automatically

You can label the partition so that it can be entered in the `fstab` file for mounting the partition automatically at bootup. The use of the label and `fstab` file is a faster operation than mounting the file system manually with a device path.

1. Type the following command to add a label to the partition:

```
e2label /dev/sdb(x) / directory-path
```

where *x* is 1 for this partition and *directory-path* is the directory that was created and its location.

2. Edit the `/etc/fstab` file and add the following line:

```
LABEL=/mount/point /mount/point ext3 1 2
```

3. Save the file.

4. To check to see if `fstab` was set up correctly, type:

```
mount -a
```

If the mount point and the `fstab` file are correctly set up, no errors appear.

5. To verify that the file system is mounted and list all mounted file systems, type:

```
df -k
```

6. To unmount the file system, type:

```
umount /filesystem-name
```

H.9 Determining the World Wide Name for Linux Hosts

Before you can create host filters, you need to know the world wide name (WWN) for the FC HBA that connects your host to your FC array.

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

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

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

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

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

See [“Setting Up Host Filter Entries”](#) on page 5-52 for more information about creating host filters.

Configuring an IBM Server Running the AIX Operating Environment

This appendix provides platform-specific host installation and configuration information to use when you connect a Sun StorEdge 3510 FC Array to an IBM server running the AIX operating environment.

The Sun StorEdge 3510 FC Array supports versions 4.3.3 and 5.1L of the IBM AIX operating environment in dual-path configurations using Sun StorEdge Traffic Manager 3.0 failover drivers for IBM AIX.

Refer to the *Sun StorEdge Traffic Manager 3.0 Installation and User's Guide for the IBM AIX Operating System* for detailed instructions about setting up the device driver on the server and for additional information about configuring your IBM server.

Customers interested in Sun StorEdge Traffic Manager 3.0 for multiplatform support should contact Sun Sales or visit:

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

For more information on multiplatform support, see:

http://www.sun.com/storage/san/multiplatform_support.html

The information in this appendix supplements the configuration steps presented in [Chapter 5](#), and covers the following steps:

- “Setting Up a Serial Port Connection” on page I-2
- “Accessing the Firmware Application From an IBM Server Running AIX” on page I-3
- “Identifying the Device on Which You Will Create a Logical Volume” on page I-4
- “Using SMIT to Enable an AIX Host to Recognize New LUNs” on page I-5
- “Creating a Volume Group” on page I-6
- “Creating a Logical Volume” on page I-7
- “Creating a File System” on page I-7

- [“Mounting the New File System”](#) on page I-8
- [“Verifying That the New File System Is Mounted”](#) on page I-9
- [“Determining the World Wide Name for IBM Servers Running AIX”](#) on page I-9

I.1 Setting Up a Serial Port Connection

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

If you are planning to access your array over an IP network or through a terminal server and only want to connect through a serial port for the initial configuration of the array, it is not necessary to configure a serial port connection from your IBM host. For convenience, installers frequently perform the initial array configuration using a serial port on a portable computer.

If you want to use a Microsoft Windows portable computer for this initial array configuration, see [“Setting Up the Serial Port Connection”](#) on page G-2 for Windows 2000 systems or [“Setting Up the Serial Port Connection”](#) on page K-2 for Windows NT systems.

If you prefer to connect through a serial port on your IBM server, consult the hardware information for your IBM host system to locate a serial port you can use for configuring the Sun StorEdge 3510 FC Array. The system documentation also tells you what device file to use in accessing that port. Then set the serial port parameters on the server. See [“Configuring a COM Port to Connect to a RAID Array”](#) on page 4-14 for the parameters to use.

Note – The next section also shows how to use the Kermit utility to set these parameters.

I.2 Accessing the Firmware Application From an IBM Server Running AIX

The RAID controller can be configured from the host system by means of terminal emulators such as Kermit.

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program after you assign an IP address to the array. For details, see [“Setting Up Out-of-Band Management Over Ethernet” on page 4-16](#) and refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

To access the controller firmware through the serial port, perform the following steps:

1. **Connect the serial null modem cable to the disk array and to the serial port on the IBM server.**

Note – A DB-9-to-DB-25 serial cable adapter is shipped with the disk array to connect the serial cable to a DB-25 serial port on your host if you do not have a DB-9 serial port.




FIGURE I-1 RAID Array COM Port Connected Locally to the Serial Port of a Host System

2. **Power on the array.**
3. **After the array is powered up, power on the IBM server and log in as root, or su to root if you are currently logged in as a user.**

4. Start the Kermit program and set the parameters as shown. Use the device-specific name for the serial port you are using. In the example, the serial port being configured is `/dev/tty0p1`.

```
# kermit
Executing /usr/share/lib/kermit/ckermi.ini for UNIX...
Good Morning!
C-Kermit 7.0.197, 8 Feb 2000, for HP-UX 11.00
  Copyright (C) 1985, 2000,
  Trustees of Columbia University in the City of New York.
Type ? or HELP for help.
(/) C-Kermit>set line /dev/tty0p1
(/) C-Kermit>set baud 38400
/dev/tty0p1, 38400 bps
(/) C-Kermit>set term byte 8
(/) C-Kermit>set carrier-watch off
(/) C-Kermit>C
Connecting to /dev/tty0p1, speed 38400.
The escape character is Ctrl-\ (ASCII 28, FS)
Type the escape character followed by C to get back,
or followed by ? to see other options.
-----
```

Note – To return to the Kermit prompt type `Ctrl \` and then `C`. To exit Kermit, first return to the Kermit prompt and then type `exit`.

I.3 Identifying the Device on Which You Will Create a Logical Volume

1. Display all the defined disk devices on your system.

```
# lsdev -Cc disk
```

A list of defined disk devices is displayed along with information about those devices.

2. Identify the name of the device on which you want to create a logical volume, for example, `hdisk7`.

3. **Verify that the disk you've identified is the one you want.**

```
# lscfg -v -l device-name
```

Detailed information about the device you specified is displayed.

4. **Examine the manufacturer information and model information to make sure you have specified the device you want.**

I.4 Using SMIT to Enable an AIX Host to Recognize New LUNs

The simplest way to configure your IBM AIX host to use your array is to use the System Management Interface Tool (SMIT). SMIT can be used with a graphical user interface (GUI) if your host is configured with a graphics card or if you want to export your display to an X-Windows terminal. SMIT can also be used with an ASCII interface that can be accessed through an ASCII terminal or console window on another system on the network. Because the ASCII interface is available on all systems, it is used as an example in the remainder of this appendix, but either interface can be used to perform the procedures described here. In addition, the same operations can be performed directly from the command line using standard AIX system administration commands.

Note – Use the Enter key when you have finished filling out a screen to advance to the next screen. Use the keystroke combinations shown on the screen to move back to previous screens.

1. **Start SMIT using the ASCII interface to display the System Management screen.**

```
# smit -a
```

2. **Choose System Storage Management (Physical and Logical Storage) from the System Management menu.**
3. **Choose Logical Volume Manager from the System Storage Management menu.**

The Logical Volume Manager menu is displayed on the Logical Volume Manager screen. You will be using menu options from this menu to create a volume group and then a logical volume within this volume group.

Logical Volume Manager
Volume Groups
Logical Volumes
Physical Volumes
Paging Space

Volume groups are a way of dividing and allocating disk storage capacity. Volume groups can be used to subdivide a large partition of storage into smaller units of usable space called *logical volumes*.

Each volume group is divided into logical volumes, which are seen by applications as individual disks. Logical volumes can contain their own file systems.

The underlying physical storage in a volume group consists of one or more physical volumes. A physical volume can be a single physical disk or a partition of a disk array. In this appendix, the physical volume is the disk device you identified in [“Identifying the Device on Which You Will Create a Logical Volume” on page I-4](#).

I.5 Creating a Volume Group

1. Choose **Volume Group** from the **Logical Volume Manager** menu.
2. Choose **Add a Volume Group** from the **Volume Groups** menu.
3. Type the name you want to give the volume group next to **VOLUME GROUP** name.
4. Next to **PHYSICAL VOLUME** name, type the name of the disk device you identified in [“Identifying the Device on Which You Will Create a Logical Volume” on page I-4](#).
An “ARE YOU SURE” confirmation screen is displayed.
5. **Confirm this message to display a status screen.**
When the volume group is created, a status screen displays “Command: OK.”
6. **Return to the Volume Groups screen to activate the new volume group.**

7. Choose **Activate a Volume Group** from the **Volume Groups** menu.
8. Next to **VOLUME GROUP** name, type the name you gave the volume group.
When the volume group is activated, a status screen displays “Command: OK.”

I.6 Creating a Logical Volume

1. Return to the **Logical Volume Manager** screen.
2. Choose **Logical Volumes** from the **Logical Volume Manager** menu to display the **Logical Volumes** screen.
3. Choose **Add a Logical Volume** from the **Logical Volumes** menu to display the **Add a Logical Volume** screen.
4. Type the name you gave the volume group next to **VOLUME GROUP** name.
The **Add a Logical Volume** screen is displayed.
5. Type the name you want to give the new logical volume next to **LOGICAL VOLUME (new)**.
6. Type the number of partitions you want to create on the new logical volume next to **Number of LOGICAL PARTITIONS**.
7. Type the values you want to use for the other parameters on the **Add a Logical Volume** screen or accept the default values.

When the logical volume is created, a status screen displays “Command: OK” and displays the name of the new logical volume.

I.7 Creating a File System

The following steps create a file system on the new logical volume.

1. Return to the **System Storage Management (Physical and Logical Storage)** screen.
The **System Storage Management** screen is displayed.
2. Choose **File Systems** from the **System Storage Management** menu to display the **File Systems** screen.
3. Choose **Add / Change / Show / Delete File Systems** from the **File Systems** menu.

4. Choose **Journalized File System** to display the **Journalized File System** screen.
5. Choose **Add a Journalized File System** from the **Journalized File System** menu.
6. Choose **Add a Standard Journalized File System** to display the **Volume Group Name** screen.
The **Volume Group Name** screen displays the names of available volume groups.
7. Move the cursor to highlight the name of the volume group you created in **“Creating a Volume Group”** on page I-6.
The **Add a Standard Journalized File System** screen is displayed.
8. Specify the size of the file system you want to create by typing the number of blocks you want next to **SIZE of file system** (in 512-byte blocks).
9. Next to **MOUNT POINT**, type the directory name (and path, if desired) on which you want to mount the new file system.

Note – If the mount point directory does not exist, it is created.

When the file system is created, a status screen displays **Command: OK**, the **Journalized File System**, the mount point, and the size of the file system.

I.8 Mounting the New File System

1. Return to the **File Systems** screen.
2. Choose **List All File Systems** from the **File Systems** menu to see the names of all file systems.
3. Identify the name of the file system you created in the **Name** column, for example, `ldev/lv00`.
4. Choose **Mount a File System** from the **File Systems** menu to display the **Mount a File System** screen.
5. Type the name of the file system you identified in Step 3 next to **FILE SYSTEM name**.
6. Type the name of the mount point you specified next to **DIRECTORY** over which to mount.
When the file system is mounted, a status screen displays **“Command: OK.”**
7. Exit **SMIT**.

I.9 Verifying That the New File System Is Mounted

1. Use the `mount` command to verify that the logical volume is mounted correctly.

```
# mount
```

The mounted file systems and directories are displayed.

2. Examine the list to make sure your new file system is mounted at the mount point you specified.
3. Display the attributes of the new file system.

```
# df -k mount-point
```

I.10 Determining the World Wide Name for IBM Servers Running AIX

Before you can create host filters, you need to know the world wide name (WWN) for the FC HBA that connects your host to your FC array.

For supported IBM HBAs, follow these steps:

1. Determine the device name by typing the command:

```
# lspcfg | grep fc
```

2. Type the following command:

```
# lspcfg -vl device-name
```

Output similar to the following is displayed. The network address is the WWN.

```

lscfg.ou.txt - Notepad
File Edit Search Help
DEVICE          LOCATION          DESCRIPTION  fcs1          20-58          FC Adapter
Part Number.....00P4494          EC Level.....A          Serial
Number.....1E3120A681          Manufacturer..... 001E          FRU
Number.....          00P4495          Network Address.....10000000C932A752          ROS Level
and ID.....02C03891          Device Specific.(Z0).....2002606D          Device
Specific.(Z1).....00000000          Device Specific.(Z2).....00000000          Device
Specific.(Z3).....02000909          Device Specific.(Z4).....FF401050          Device
Specific.(Z5).....02C03891          Device Specific.(Z6).....06433891          Device
Specific.(Z7).....07433891          Device Specific.(Z8).....20000000C932A752          Device
Specific.(Z9).....CS3.82A1          Device Specific.(ZA).....C1D3.82A1          Device
Specific.(ZB).....C2D3.82A1          Device Specific.(ZL).....P1-I1/Q1
  
```

FIGURE I-2 Network Address Corresponding to WWN

Configuring an HP Server Running the HP-UX Operating Environment

This appendix provides platform-specific host installation and configuration information to use when you connect a Sun StorEdge 3510 FC Array to an HP server running the HP-UX operating environment.

For a list of supported host bus adapters, refer to the *Sun StorEdge 3000 Family Release Notes* for your array.

The Sun StorEdge 3510 FC Array supports the HP-UX operating environment Level 11.0 and Level 11.i in dual-path configurations using Sun StorEdge Traffic Manager 3.0 failover drivers for the HP-UX operating environment.

Refer to the *Sun StorEdge Traffic Manager 3.0 Installation and User's Guide for the Hewlett Packard HP-UX Operating System* for detailed instructions about setting up the device driver on the server and for additional information about configuring your HP server.

Customers interested in Sun StorEdge Traffic Manager 3.0 for multiplatform support should contact Sun Sales or visit:

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

For more information on multiplatform support, see:

http://www.sun.com/storage/san/multiplatform_support.html

The information in this appendix supplements the configuration steps presented in [Chapter 5](#), and covers the following steps:

- [“Setting Up a Serial Port Connection”](#) on page J-2
- [“Accessing the Firmware Application From an HP Server Running HP-UX”](#) on page J-3
- [“Attaching the Disk Array”](#) on page J-5

- [“Logical Volume Manager” on page J-6](#)
- [“Definitions of Common Terms” on page J-6](#)
- [“Creating a Physical Volume” on page J-7](#)
- [“Creating a Volume Group” on page J-7](#)
- [“Creating a Logical Volume” on page J-9](#)
- [“Creating an HP-UX File System” on page J-10](#)
- [“Mounting the File System Manually” on page J-10](#)
- [“Mounting the File System Automatically” on page J-10](#)
- [“Determining the World Wide Name for HP-UX Hosts” on page J-11](#)

J.1 Setting Up a Serial Port Connection

The RAID controller can be configured by means of a host system running a VT100 terminal emulation program or by a Microsoft Windows terminal emulation program such as Hyperterminal.

If you are planning to access your array over an IP network or through a terminal server and only want to connect through a serial port for the initial configuration of the array, it is not necessary to configure a serial port connection from your HP host. For convenience, installers frequently perform the initial array configuration using a serial port on a portable computer.

If you want to use a Microsoft Windows portable computer for this initial array configuration, see [“Setting Up the Serial Port Connection” on page G-2](#) for Windows 2000 systems or [“Setting Up the Serial Port Connection” on page K-2](#) for Windows NT systems.

If you prefer to connect through a serial port on your HP server, consult the hardware information for your HP host system to locate a serial port you can use for configuring the Sun StorEdge disk array. The system documentation also tells you what device file to use in accessing that port. Then set the serial port parameters on the server. See [“Configuring a COM Port to Connect to a RAID Array” on page 4-14](#) for the parameters to use.

Note – The next section also shows how to use the Kermit utility to set these parameters.

Once you have configured your serial port, follow the instructions in the next section.

J.2 Accessing the Firmware Application From an HP Server Running HP-UX

The RAID controller can be configured from the host system by means of terminal emulators such as `cu` or Kermit. These instructions show the use of Kermit. For information on `cu` see `cu(1)`.

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program after you assign an IP address to the array. For details, see [“Setting Up Out-of-Band Management Over Ethernet” on page 4-16](#) and refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

To access the controller firmware through the serial port, perform the following steps:

1. Use a null modem serial cable to connect the COM port of the RAID array to an unused serial port on your host system.

A null modem cable has serial signals swapped for connecting to a standard serial interface.

Note – A DB-9-to-DB-25 serial cable adapter is included in your package contents to connect the serial cable to a DB-25 serial port on your host if you do not have a DB-9 serial port.




FIGURE J-1 RAID Array COM Port Connected Locally to the Serial Port of a Host System

2. Power on the array.
3. After the array is powered up, power on the HP server and log in as root, or su to root if you are logged in as a user.
4. Start the Kermit program and set the parameters as shown. Use the device-specific name for the serial port you are using. In the example, the serial port being configured is /dev/tty0p1.

```
# kermit

Executing /usr/share/lib/kermit/ckermid.ini for UNIX...

Good Morning!

C-Kermit 7.0.197, 8 Feb 2000, for HP-UX 11.00

Copyright (C) 1985, 2000,

Trustees of Columbia University in the City of New York.

Type ? or HELP for help.

(/) C-Kermit>set line /dev/tty0p1

(/) C-Kermit>set baud 38400

/dev/tty0p1, 38400 bps

(/) C-Kermit>set term byte 8

(/) C-Kermit>set carrier-watch off

(/) C-Kermit>C

Connecting to /dev/tty0p1, speed 38400.

The escape character is Ctrl-\ (ASCII 28, FS)

Type the escape character followed by C to get back,

or followed by ? to see other options.
```

Note – To return to the Kermit prompt, type `Ctrl \` and then `C`. To exit Kermit, first return to the Kermit prompt and then type `exit`.

J.3 Attaching the Disk Array

The simplest way to configure a disk array is to use System Administration Manager (SAM), HP-UX's system administration tool. If SAM is not installed on your system, or if you prefer to use the command-line interface, the following procedures guide you through the task. For more information please consult the HP document, *Configuring HP-UX for Peripherals*:

1. Use the `ioscan` command to determine what addresses are available on the HBA to which you will be attaching the Sun StorEdge 3510 FC Array.
2. Access the firmware application on the Sun StorEdge 3510 FC Array and set the SCSI IDs of the host channels you will be using.
3. Map the partitions containing storage that you want to use to the appropriate host channels.
Partitions must be assigned to LUNs in sequential order, beginning at LUN 0.
4. Halt the operating system using the `shutdown` command.
5. Turn off all power to peripheral devices and then to the server.
6. Attach one or more host channels of the Sun StorEdge 3510 FC Array to the SCSI interface cards in the host using the supplied SCSI cables.
7. Turn on the power to the Sun StorEdge 3510 disk array and all other peripheral devices. After they are initialized, power on the server and boot HP-UX. During the boot process, the operating system recognizes the new disk devices and builds device files for them.
8. Verify that you can see the new storage resources by running the `ioscan` command. You are now ready to use the storage.

Note – If you create and map new partitions to the array, you can have them recognized by the operating system without rebooting. Run the `ioscan` and the `insf` commands to discover the resources and to create their device files.

J.4 Logical Volume Manager

The Logical Volume Manager (LVM) is a disk management system provided by HP in all versions of HP-UX 11. The LVM allows you to manage storage as logical volumes. This section describes some concepts used by the LVM and explains how to create logical volumes on your Sun StorEdge 3510 FC Array. For more detailed information on the LVM, please consult `lvmd(7)` and the HP publication *Managing Systems and Workgroups: Guide for HP-UX System Administration* (HP part number B2355-90742).

As with many system administration tasks, you can use SAM to create and maintain logical volumes. However, some functions can only be performed with HP-UX commands. The procedures in this appendix are performed using the command-line interface rather than SAM.

J.5 Definitions of Common Terms

Volume groups are HP-UX's method for dividing and allocating disk storage capacity. Volume groups can be used to subdivide a large partition of storage into smaller units of usable space called *logical volumes*.

Each volume group is divided into logical volumes, which are seen by the applications as individual disks. They can be accessed as either character or block devices and can contain their own file systems.

The underlying physical storage in a volume group consists of one or more *physical volumes*. A physical volume can be a single physical disk or a partition of a disk array.

Each physical volume is divided into units called *physical extents*; the default size of these units is 4 MB, but can range in size from 1 MB to 256 MB. The maximum number of physical extents that a volume group can contain is 65,535. With the default size of 4 MB, this limits the size of the volume group to 255 GB.

To create a volume group larger than 255 GB, you must increase the size of the physical extents when creating the volume group. See `vgcreate(1m)` for further information.

J.6 Creating a Physical Volume

To use a storage resource in the LVM, it must first be initialized into a physical volume (also called an *LVM disk*).

Log in as root, or su to root if you are not logged in with root user privileges.

1. **Select one or more partitions on the Sun StorEdge 3510 FC array that you want to use. The output of `ioscan(1M)` shows the disks attached to the system and their device names:**

```
# ioscan -fnC disk

Class I   H/W Path      Driver  S/W State  H/W Type  Description
-----
disk  1   0/12/0/0.6.0 sdisk   CLAIMED   DEVICE    Sun StorEdge 3510

/dev/dsk/c12t6d2 /dev/rdisk/c12t6d2
```

2. **Initialize each partition as an LVM disk with the `pvcreate` command. For example, type:**

```
# pvcreate /dev/rdisk/c12t6d2
```



Caution – This process results in the loss of any data that resides on the partition.

J.7 Creating a Volume Group

The volume group contains the physical resources that you can use to create usable storage resources for your applications.

1. Create a directory for the volume group and a device file for the group in that directory:

```
# mkdir /dev/vgmynewvg  
  
# mknod /dev/vgmynewvg/group c 64 0x060000
```

The name of the directory is the name of the volume group. By default, HP-UX uses names of the format *vgNN*, but you can choose any name that is unique within the list of volume groups.

In the preceding example, the `mknod` command has the following arguments:

- The fully qualified path name of the new device file `group`
- The letter `c` indicates a character device file
- The major number `64` (used for all volume groups)
- A minor number of the form `0xNN0000`, where *NN* is the two-digit hexadecimal representation of the volume group number (`06` in the example)

To associate the physical volume with a volume group, use the `vgcreate` command:

```
# vgcreate /dev/vgmynewvg /dev/dsk/c12t6d2
```

To verify the creation and view the volume group properties, use the `vgdisplay` command:

```
# vgdisplay vg02  
--- Volumegroups ---  
  
VG Name                /dev/vg02  
VG Write Access        read/write  
VG Status               available  
Max LV                 255  
Cur LV                 0  
Open LV                0  
Max PV                 16  
Cur PV                 1  
Act PV                 1  
Max PE per PV          2167  
VGDA                   2
```

PE Size (Mbytes)	4
Total PE	2167
Alloc PE	0
Free PE	2167
Total PVG	0

In the output of `vgdisplay`, the Total PE field gives the number of physical extents in the volume group.

The size of each physical extent is given in the PE Size field (the default is 4 MB), so the total capacity of this volume group is $2167 \times 4 \text{ MB} = 8668 \text{ MB}$.

The Alloc PE field shows the number of physical extents allocated to logical volumes. At this point, the Alloc PE field is zero because we have not assigned any of this volume group's capacity to logical volumes.

J.8 Creating a Logical Volume

To create a logical volume within the volume group, use the `lvcreate` command with the `-L` option to specify the size of the logical volume in megabytes. The logical volume size should be a multiple of the physical extent size. In this example, a logical volume of 4092 MB is created:

```
# lvcreate -L 4092 /dev/vg02
```

Both character and block device files for the new logical volume are created in the volume group directory:

```
# ls /dev/vg02

group   lv011   rlv011
```

Applications should use these names to access the logical volumes. Unless you specify otherwise, HP-UX creates names in the form shown in the example. To specify custom names for logical volumes see `vgcreate(1M)`.

J.9 Creating an HP-UX File System

The following command creates a file system on the logical volume created in the previous steps.

```
# /sbin/newfs -F vxfs /dev/vgmynewvg/rlvol1
```

J.10 Mounting the File System Manually

The process of incorporating a file system into the existing directory structure is known as “mounting the file system.” The files, although present on the disk, are not accessible to users until they are mounted.

1. **Create a directory to be the mount point for your new file system:**

```
# mkdir /usr/local/myfs
```

2. **To mount your file system, type the following:**

```
# mount /dev/vgmynewvg/lvol1 /usr/local/myfs
```

J.11 Mounting the File System Automatically

By placing information about your file system in the `fstab` file, you can have HP-UX mount the file system automatically during bootup. You can also use the name of the mount point in mount commands that you issue from the console.

1. **Make a copy of the existing `fstab` file:**

```
# cp /etc/fstab /etc/fstab.orig
```

2. To include the file system created in the example, add the following line to the file `/etc/fstab`:

```
/dev/vg0mynewvg/lvol1 /usr/local/myfs vxfs delaylog 0 2
```

See the entry for `fstab(4)` for details on creating `/etc/fstab` entries.

3. To check to see if `fstab` was set up correctly, type:

```
# mount -a
```

If the mount point and the `fstab` file are correctly set up, no errors appear.

4. To verify that the file system is mounted and list all mounted file systems, type:

```
# bdf
```

5. To unmount the file system, type:

```
# umount /usr/local/myfs
```

J.12 Determining the World Wide Name for HP-UX Hosts

Before you can create host filters, you need to know the world wide name (WWN) for the FC HBA that connects your host to your FC array.

For supported HP-UX host HBAs, follow these steps:

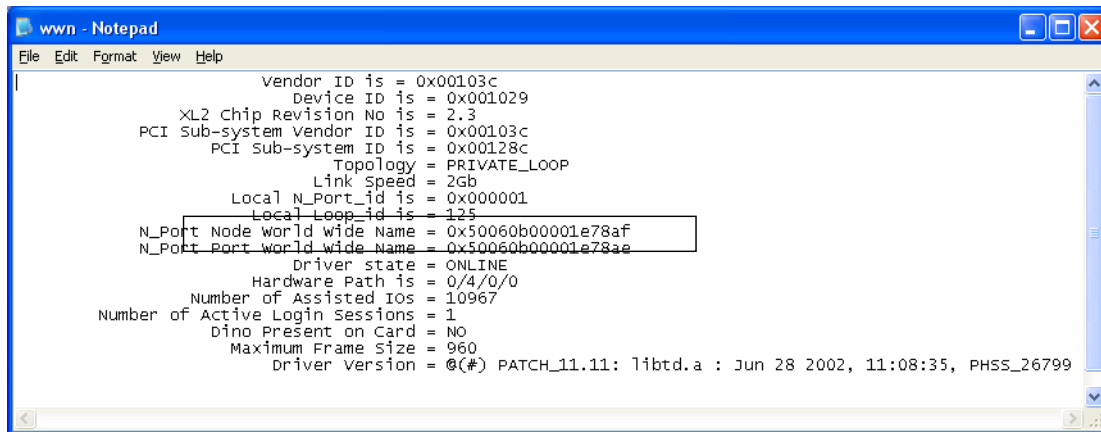
1. Determine the device name by typing the command:

```
# ioscan -fnC fc
```

2. Type:

```
# fcmsutil/device-name/
```

Output similar to the following is displayed:



```
Vendor ID is = 0x00103c
Device ID is = 0x001029
XL2 Chip Revision No is = 2.3
PCI Sub-system Vendor ID is = 0x00103c
PCI Sub-system ID is = 0x00128c
Topology = PRIVATE_LOOP
Link speed = 2Gb
Local N_Port_id is = 0x000001
Local Loop_id is = 125
N_Port Node world wide Name = 0x50060b00001e78af
N_Port Port world wide Name = 0x50060b00001e78ae
Driver state = ONLINE
Hardware Path is = 0/4/0/0
Number of Assisted IOS = 10967
Number of Active Login Sessions = 1
Dino Present on Card = NO
Maximum Frame Size = 960
Driver version = @(#) PATCH_11.11: libtd.a : Jun 28 2002, 11:08:35, PHSS_26799
```

The Node World Wide Name shown is the WWN you use when configuring the RAID controller.

Configuring a Windows NT Server

This appendix provides platform-specific host installation and configuration information to use when you connect a Sun StorEdge 3510 FC Array to a host running the Windows NT Server operating system.

The Sun StorEdge 3510 FC Array supports Microsoft Windows NT 4.0, Service Pack 6A or greater, in either single-path or dual-path configurations. Dual-path configurations use the Sun StorEdge Traffic Manager 3.0 for Microsoft Windows NT.

If you are going to use dual-path instead of single-path configurations, see the *Sun StorEdge Traffic Manager 3.0 Installation and User's Guide for Microsoft Windows 2000 and Windows NT* for detailed instructions about setting up the device driver on the server and for additional information about configuring your Windows NT server.

Customers interested in Sun StorEdge Traffic Manager 3.0 for multiplatform support should contact Sun Sales or visit:

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

For more information on multiplatform support, see:

http://www.sun.com/storage/san/multiplatform_support.html

The information in this appendix supplements the configuration steps presented in [Chapter 5](#), and covers the following steps:

- [“Setting Up the Serial Port Connection” on page K-2](#)
- [“Accessing the Firmware Application From a Windows NT Server” on page K-5](#)
- [“Enabling a Windows NT Server to Recognize New Devices and LUNs” on page K-5](#)
- [“Determining the World Wide Name for Windows NT Servers” on page K-9](#)

K.1 Setting Up the Serial Port Connection

The RAID controller can be configured by means of a host system running a VT100 terminal emulation program or by a Microsoft Windows terminal emulation program such as Hyperterminal.

Note – You can also monitor and configure a RAID array over an IP network with the Sun StorEdge Configuration Service program after you assign an IP address to the array. For details, see [“Setting Up Out-of-Band Management Over Ethernet” on page 4-16](#) and refer to the *Sun StorEdge 3000 Family Configuration Service User’s Guide*.

1. **Use a null modem serial cable to connect the COM port of the RAID array to an unused serial port on your host system.**

Note – A DB-9-to-DB-25 serial cable adapter is included in your package contents to connect the serial cable to a DB-25 serial port on your host if you do not have a DB-9 serial port.

2. **Power up the array.**
3. **On the server, choose Start → Programs → Accessories → Hyperterminal → HyperTerminal.**
4. **Type a name and choose an icon for the connection.**

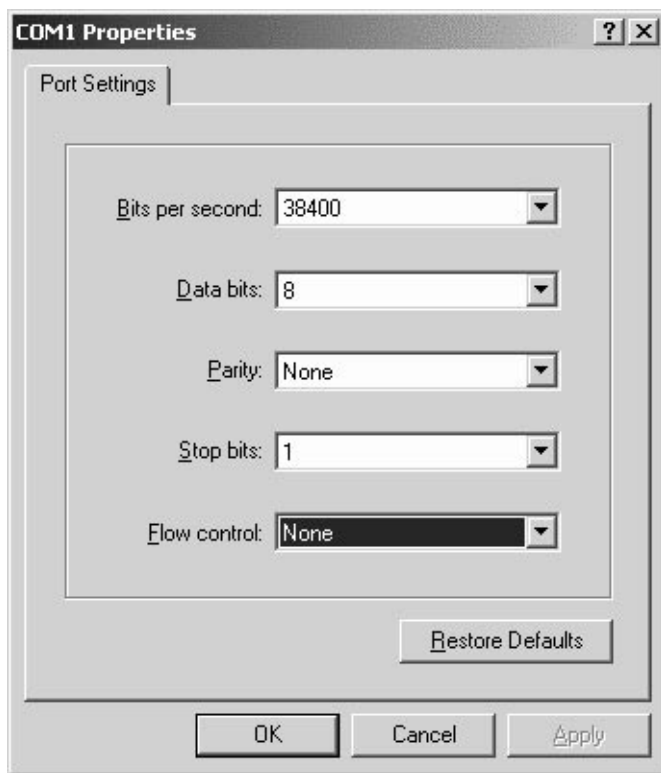
5. In the Connect To window, select the COM port from the “Connect using:” drop-down menu that is connected to the array.



6. Click OK.

7. In the Properties window, set the serial port parameters using the drop-down menus.

See “Configuring a COM Port to Connect to a RAID Array” on page 4-14 for the parameters to use.



8. Click OK.

9. To save the connection and its settings, choose File → Save.

The connection file name is *connection-name*, where *connection-name* is the name you gave this HyperTerminal connection when you created it.

10. To make a connection shortcut on your desktop, choose Start → Find → For Files or Folders. Enter the *connection-name* and click the Search Now button. Highlight and right-click on the file name in the Search Results window, choose Create Shortcut, and click Yes.

You are now ready to configure your array.

K.2 Accessing the Firmware Application From a Windows NT Server

To access the array from a Windows NT server through a serial port, use the HyperTerminal connection you configured in [“Setting Up the Serial Port Connection” on page K-2](#) or install a VT100 terminal emulation program and use the parameters described there.

To access the array through an Ethernet connection, follow the procedures in [“Setting an IP Address” on page 4-14](#). To access the command prompt described there from a Windows NT server, choose Programs → Command Prompt.

K.3 Enabling a Windows NT Server to Recognize New Devices and LUNs

Before beginning this procedure, make sure that your system has been upgraded with Service Pack 2 or higher.

Before beginning this procedure, make sure that you are using a supported FC HBA such as a QLogic QLA2310 or QLogic QLA2342. Refer to the release notes for your array for current information about which HBAs are supported.

Also make sure that you are using a supported driver for your HBA. For the QLA2310 or QLA2342, use driver version 8.1.5.12 or later.

Note – When the Microsoft Windows operating system recognizes the Sun StorEdge 3510 FC Array’s SCSI Enclosure Services (SES) processor and identifies it as a SE3510 Scsi Enclosure, it might prompt you to supply a driver. This is standard plug-and-play behavior when unknown devices are encountered, but no device driver is necessary. If you see this prompt, simply press Cancel. If you have multiple arrays, you might see this prompt more than once. Each time you see this prompt, press Cancel.

1. **Boot your system and verify that the HBA basic input/output system (BIOS) recognizes your new FC device.**

Note – While your system is starting up, you should see your new Fibre Channel devices.

You are now ready to partition and format your new devices.

2. Open the Disk Administrator application.

a. Click Start.


b. Choose Administrative Tools from the Programs menu.

c. Choose Disk Administrator.

A “Disk Administrator is initializing” progress indicator is displayed.

A Disk Administrator window then displays the drives recognized by the system.

3. Select the disk whose free space you want to partition and format.



a. Choose Create from the Partition menu.


A Create Primary Partition dialog box enables you to specify the size of the partition.

b. Specify a size or accept the default.

c. Click OK to create the partition.

The partition is now identified as Unformatted in the Disk Administrator window.

4. Select the Unformatted partition.



5. Choose Commit Changes Now from the Partition menu.

A confirmation dialog box is displayed.

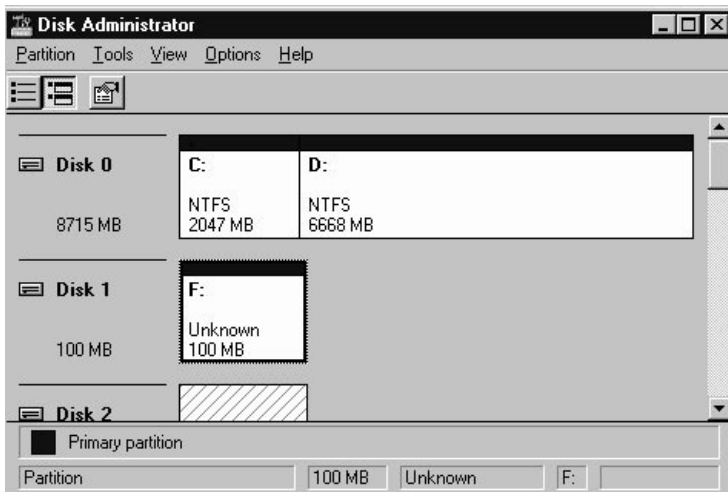
6. Click Yes to save the changes you have made.

A dialog box confirms that disks were updated successfully.

7. Click OK.

The partition is now identified as Unknown in the Disk Administrator window.

8. Format the Unknown partition.



a. Select the Unknown partition.

b. Choose Format from the Tools menu.

A Format dialog box is displayed.

c. Choose NTFS from the File System drop-down menu.

d. Make sure the Quick Format checkbox is selected.

e. When you have specified the settings you want, click Start.

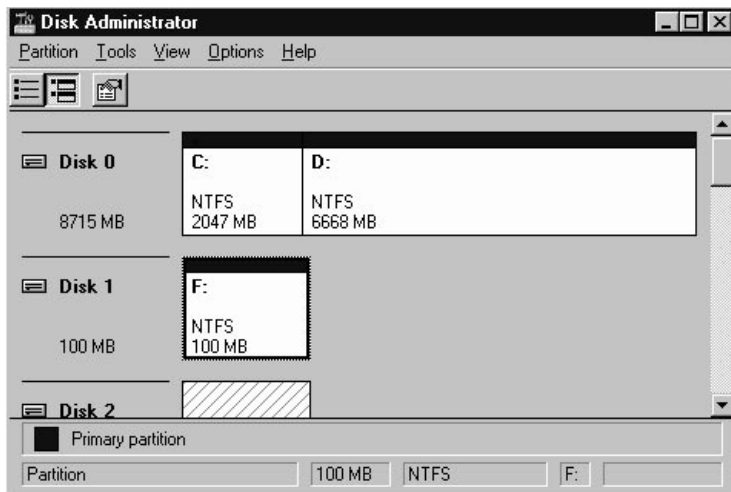
A dialog box warns you that any existing data on the disk will be erased.

f. Click OK to format the disk.

The new partition is formatted and a dialog box confirms that the format is complete.

9. Click OK.

The formatted partition is identified as NTFS in the Disk Administrator dialog box.



10. Repeat these steps for any other new partitions and devices you want to format.

K.4 Determining the World Wide Name for Windows NT Servers

Before you can create host filters, you need to know the world wide name (WWN) for the FC HBA that connects your host to your FC array.

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

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

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

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

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

See [“Setting Up Host Filter Entries”](#) on page 5-52 for more information about creating host filters.

Index

SYMBOLS

#FL, 7-10
#LN, 7-10
#SB, 7-10

A

active-active configuration, 5-2
active-standby configuration, 5-2
agent
 Agent Options Management window, B-13
Agent Options Management window, B-5
alarms
 failed component description, C-1
 silencing, 7-5
amber blinking on Event LED, 6-4
array
 connecting, 4-1
 mounting, 3-5
automatic rebuild, 8-3

B

BAD drive status, 7-12
battery, 5-4
 date code, 5-4
 shelf life, 5-4
 when to change, 5-4
Battery Operation, 5-4
baud rate, 4-14

beep codes
 muting, 7-5
bus adapter
 connecting to, 4-23

C

cables
 customer-provided, 3-5
cabling
 cables supplied, 3-3
 Ethernet connection, 4-16
 procedures, 4-1
 RS-232 connection, 4-14
Cabling to Expansion Units, 4-17, B-9
cache options, 5-6
cache status, 7-2
CD supplied, 1-11
channels
 default channels, 5-22
 drive, 4-10
 host, 4-10
 host or drive, 5-22
checking package contents, 3-3
Checking Status Windows, 7-8
clearances to allow, 2-5
CLI. See Sun StorEdge CLI.
COM port
 connecting to, 4-14
COM port connection, F-2, G-2, J-3
COM port parameters, F-2, G-4, H-2, I-2, I-4, J-2, J-5

- communication parameters, 4-14
- concurrent rebuild, 8-6
- configuration
 - key steps, 5-7
 - minimum requirements, 5-31
 - overview, 5-7
 - restoring from a file, 7-24
 - serial port connection, F-1, G-2, J-2, K-2
- Configuring a Linux Server, H-1
- Configuring a Sun Server Running the Solaris Operating Environment, F-1
- Configuring a Windows 2000 Server or Windows 2000 Advanced Server, G-1
- Configuring a Windows NT Server, B-1, C-1, K-1
- connecting an expansion unit, 4-17, B-9
- Connecting Ports to Hosts, 4-23
- Connecting SFPs to SFP Ports, 4-13, B-12
- Connecting the Chassis to an AC Power Outlet, 4-6
- Connecting the Chassis to DC Power Outlets, 4-7
- Connecting Your Fibre Channel Array, 4-1
- controller assignment, 5-39
- Controller Failover, 8-2
- controller failure, 8-2
- controller limitations, 5-2
- controller NAME, 5-42
- controller voltage and temperature, 7-15
- create
 - logical drive, 5-32
- customer obligations, 2-2
- customer-provided cables, 3-5

D

- Deleting Logical Drives, 7-6
- DHCP, 4-15
- disks not visible, 8-1
- drive
 - identifying a failed, 8-7
- drive channels, 4-10
- drive failure
 - recovering from a fatal, 8-11
- DRV FAILED status, 7-10
- dual-bus configurations, B-6
- Dynamic Host Configuration Protocol, 4-14

E

- electrical specifications, 2-4
- electromagnetic compatibility (EMC), 2-3
- environmental range, 2-3
- Environmental Requirements, 2-3
- Ethernet connection, 4-16
- event LED, 6-3, B-19
- Event Logs, 7-22
- expansion unit
 - setting loop ID, 4-20, B-7
- expansion units
 - cabling to, 4-17, B-9

F

- failed drive
 - identifying, 8-7
- failover
 - controller, 8-2
- fan LED, 6-3, B-19
- fatal drive failure
 - recovering from, 8-11
- FC Protocols, 1-6
- FC Topologies, 1-7
- Fibre Channel
 - loop architecture, 1-8
 - protocols, 1-6
 - technology overview, 1-6
 - topologies, 1-7
- Field-Replaceable Units, 1-3
- filter entries
 - host, 5-52
- firmware
 - cache status, 7-2
 - controller failure, 8-2
 - controller voltage and temperature, 7-15
 - cursor bar, 7-2
 - gauge range, 7-2
 - initial window, 5-19, 7-2
 - logical drive status, 7-9
 - main menu, 7-4
 - manual rebuild, 8-5
 - navigation keys, 5-19
 - quick installation, 7-4
 - SCSI channel status, 7-13

- SCSI drive status, 7-11
- transfer rate indicator, 7-2
- upgrading, 7-26
- view event logs, 7-22
- VT-100 screen selection, 7-2
- VT100 terminal mode, 7-2
- Flash All Drives command, 8-8
- flash drive time command, 8-8
- flashing
 - all SCSI drives, 8-10
 - selected SCSI drives, 8-9
- format command, 5-45
- Front-Panel LEDs, 6-1, B-17
- FRUs, 1-3

G

- gauge range, 7-2

H

- hardware PLD code version conflict, 6-4
- hardware specifications, A-3
- Host and Drive Channels, 4-10
- host bus adapter
 - connecting to, 4-23
- host channels, 4-10
- host configuration
 - Linux, H-1
 - Solaris operating environment, F-1
 - Windows 2000, G-1
 - Windows NT, B-1, C-1, K-1
- host connections, 4-23
- host filter, 5-52
- Host Filter Entries, 5-52
- host IDs, adding, 5-25
- host IDs, creating, 5-25
- host LUNs
 - not visible, 8-1
- humidity, 2-3

I

- ID switch, 4-20, B-7
- Identifying a Failed Drive for Replacement, 8-7
- identifying drives, 8-7
- Identifying SCSI drive command, 8-8
- in-band connections, 5-6
- INCOMPLETE status, 7-10
- initial firmware window, 5-19
- INITING status, 7-9
- INVALID status, 7-9
- IP address, setting, 4-16

J

- JBOD
 - dual-bus configuration, B-6, B-14
 - single-bus configuration, B-5, B-14

K

- Key Screens and Commands, 7-2
- keys to navigate, 5-19

L

- layout map, 2-5
- LEDs
 - drives, 6-1, B-17
 - front panel, 6-1, 6-3, B-17, B-19
 - rear panel, 6-5, B-19
- Level 3
 - NEBS, 1-1
- limitations
 - controller, 5-2
- local spare assignment, 5-35
- logical drive
 - 128-LUN limit, 5-29
 - assign local spare, 5-35
 - assign RAID level, 5-33
 - change assignment, 5-39
 - creating, 5-31, 5-32
 - defaults, 5-29
 - ID, 7-9
 - larger than 253 Gbyte, 5-38

- LG number, 7-9
- maximum number of disks per logical drive, 5-28
- maximum physical capacity, 5-35
- maximum physical drive capacity, 5-35
- maximum usable capacity per logical drive, 5-29
- minimum requirements, 5-31
- RAID level, 7-9
- RAID level selected, 5-33
- RAID levels defined, 5-30
- size, 7-9
- status, 7-9
- status table, 7-9
- logical drive preferences, 5-35
- logical drives
 - deleting, 7-6
 - rebuilding, 8-3
- logical volume, 5-30
- loop id
 - setting on expansion units, 4-20, B-7
- LUN
 - described, 5-46

M

- main menu, 7-4
- management tools
 - accessing, 5-6
- manual rebuild, 8-5
- map
 - layout, 2-5
- maximum drive capacity, 5-35
- MISSING drive status, 7-12
- Mounting Your FC Array, 3-5
- muting
 - alarms, 7-5
 - beeper, 7-5

N

- NAME (controller), 5-42
- NEBS
 - Level 3, 1-1
- Network Equipment Building System, 1-1
- NEW DRV drive status, 7-11

- null modem cable, 4-14
- null modem serial cable, 4-14
- NVRAM
 - restoring, 7-24

O

- optimization
 - random I/O
 - maximum size, 5-28
 - sequential I/O
 - maximum size, 5-28
- out-of-band connections, 5-6
- out-of-band management, 4-16

P

- package contents, 3-3
 - checking, 3-3
- parameters
 - serial port, 4-14
- physical drive
 - size and speed, 7-11
 - status, 7-11
- physical drive capacity setting, 5-35
- planning site, 2-1
- ports
 - connecting (see cabling), 4-1
- power LED, 6-3, B-19
- power outlets
 - connecting to AC, 4-6
 - connecting to DC, 4-7
- power specifications, 2-4
- powering up an array, 4-9
- Powering Up and Checking LEDs, 4-9
- preinstallation worksheet, 2-7
- preparing for logical drives larger than 253 Gbyte, 5-38
- probe-scsi-all command, 5-45
- Product and Architecture Overview, 1-1
- protocols
 - Fibre Channel, 1-6

Q

quick installation (not applicable), 7-4

R

rack-mounting an array

 cabinet-mounting an array, 3-5

RAID (1+0)

 concurrent rebuild in, 8-6

RAID (5+0), 5-30

RAID level assignment, 5-33

RAID levels

 defined, 5-30

 selected, 5-33

RAID levels defined, 5-30

random I/O optimization

 maximum size, 5-28

RARP, 4-15

Rear-Panel LEDs, 6-5, B-19

rebuild

 automatic, 8-3

 concurrent in RAID (1+0), 8-6

 manual, 8-5

Rebuilding Logical Drives, 8-3

Recovering From Fatal Drive Failure, 8-11

redundant controller

 explained, 8-2

release notes, 1-11

requirements

 environmental, 2-3

reset button, 6-2, 7-6

 silencing failed component alarms, 6-2, 7-5

Restoring Your Configuration (NVRAM) From a File, 7-24

RS-232

 connecting to, 4-14, G-2, J-3

 upgrading firmware via, 7-28

S

safety precautions, 2-2

SB-MISS drive status, 7-12

screens and commands

 key, 7-2

SCSI channel status, 7-13

SCSI drive

 global or local spare, 7-11

 STANDBY mode, 7-11

 USED DRV mode, 7-11

 vendor ID, 7-12

SCSI drive status, 7-11

sequential I/O optimization

 maximum size, 5-28

serial cable

 null modem, 4-14

serial port

 parameters, 4-14

serial port connection, F-1, G-2, J-2, K-2

serial port parameters, F-2, G-4, H-2, I-2, I-4, J-2, J-5

SES version conflict, 6-4

Setting Loop IDs on Expansion Units, 4-20, B-7

settings records to keep, D-1

SFP connectors, 4-23

 inserting in ports, 4-13, B-12

silencing alarms, 7-5

 reset button, 6-2

site planning, 2-1

 console requirement, 2-7

 customer obligations, 2-2

 electrical/power, 2-4

 EMC, 2-3

 environmental, 2-3

 layout, 2-5

 physical, 2-5

 safety precautions, 2-2

site preparation survey, 2-7

size of drive, 7-11

software

 accessing the array, 5-6

 installing, 5-59

software tools, 1-11, 5-59

spare (local for logical drive), 5-35

spare drives, 8-6

specifications

 clearances, 2-5

 electrical/power, 2-4

 physical (array), 2-5

specifications (product), A-3

speed of drive, 7-11

STAND-BY drive status, 7-11

- status windows
 - checking, 7-8
- Sun StorEdge CLI, B-7
- Sun StorEdge Configuration Service, B-5
- switch
 - ID, 4-20, B-7

T

- temp LED, 6-3, B-19
- temperature
 - environmental range, 2-3
- temperature LED, 6-3, B-19
- topologies
 - Fibre Channel, 1-7
- transfer rate indicator, 7-2
- troubleshooting
 - disks not visible, 8-1
 - LUNs not visible, 8-1

U

- Upgrading Firmware, 7-26
- USED DRV drive status, 7-11

V

- VERITAS DMP, enabling, 5-61
- Veritas software, 5-61
- VT100 terminal connection, 4-14
- VT100 terminal mode, 7-2

W

- worksheet
 - preinstallation, 2-7
- worldwide name
 - determining on Solaris OE, F-3
- write-back cache, 5-6
- write-through cache, 5-6