



Sun StorEdge™ A1000 and A3x00/A3500FC Best Practices Guide

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

Part No. 806-6419-14
[November 2002, Revision A](#)

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

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

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this 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 in other countries.

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

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

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

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

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

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

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

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

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, Ultra, Ultra Enterprise, RSM, SunSolve, 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.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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



Contents

Preface **xiii**

1. Troubleshooting Overview **1-1**

1.1 A3x00/A3500FC Commandments 1-2

1.2 Available Tools and Information 1-3

 1.2.1 Documentation 1-3

 1.2.2 Web Sites 1-4

 1.2.3 Internal Directory 1-4

 1.2.4 Obtaining the Latest Version of RAID Manager 1-4

 1.2.5 RAID Manager 6.0, 6.1 and 6.22 Are not Supported 1-5

 1.2.6 Serial Cable 1-5

 1.2.7 RAID Manager 6.xx Architecture White Paper Available 1-6

1.3 Tips for Filing a Bug 1-6

1.4 FINs and FCOs 1-6

2. Hardware Installation and Configuration **2-1**

2.1 New Installation 2-2

 2.1.1 Battery Unit 2-2

 2.1.2 Power Cables 2-2

 2.1.3 Power Sequencer 2-3

- 2.1.4 Local/Remote Switch 2-3
- 2.1.5 SCSI and Fiber-Optic Cables 2-3
- 2.1.6 SCSI ID, Loop ID, Controller, and Disk Tray Switch Settings 2-4
- 2.1.7 World Wide Name (WWN) 2-5
- 2.2 Adding or Moving Arrays to a Host With Existing A3x00 Arrays 2-6
- 2.3 Adding Disks or Disk Trays 2-6
 - 2.3.1 Adding or Moving Disk Trays to Existing Arrays 2-7
 - 2.3.2 Adding or Moving Disk Drives to Existing Arrays 2-7
- 2.4 Setting Up 2x7 and 3x15 Configurations, and Converting 1x5 to 2x7 or 3x15 2-8
- 2.5 Sun StorEdge A3500/A3500FC Lite 2-9
- 2.6 Cluster, Multi-Initiator, and SCSI Daisy Chaining Configurations 2-10
 - 2.6.1 Cluster Information 2-10
 - 2.6.2 Multi-Initiator Information 2-11
 - 2.6.3 SCSI Daisy Chaining Information 2-11
- 2.7 Supported Configurations 2-11
 - 2.7.1 Maximum Server Configurations 2-12
 - 2.7.2 Onboard SOC+ 2-12
 - 2.7.3 Second Port on the SOC+ Card 2-13
 - 2.7.4 Disk Drive Support Matrices 2-13
 - 2.7.5 Independent Controller/Box Sharing 2-13
 - 2.7.6 HBAs 2-13
- 2.8 SCSI to FC-AL Upgrade 2-14

3. RAID Manager Installation and Configuration 3-1

- 3.1 Installation and Configuration Tips, Tunable Parameters, and Settings 3-2
 - 3.1.1 Software Installation 3-2
 - 3.1.2 Software Configuration 3-3
 - 3.1.3 RAID Module Configuration 3-3

3.1.4	Tunable Parameters and Settings	3-3
3.1.5	Multi-Initiator/Clustering Environment	3-4
3.1.6	Maximum LUN Support	3-5
3.2	LUN Creation/RAID Level	3-5
3.2.1	General Information	3-5
3.2.2	LUN Numbers	3-6
3.2.3	The Use of RAID Levels	3-6
3.2.4	Cache Mirroring	3-6
3.2.5	Reconstruction Rate	3-7
3.2.6	Creation Process (Serial/Parallel) Time	3-8
3.2.7	DacStor Size (Upgrades)	3-8
3.3	LUN Deletion and Modification	3-9
3.4	Controller and Other Settings	3-9
3.4.1	NVSRAM Settings	3-9
3.4.2	Parity Check Settings	3-10
3.4.2.1	RAID Manager 6.1.1	3-10
3.4.2.2	RAID Manager 6.22x	3-10
3.4.2.3	Parity Repair	3-11
3.4.2.4	Multi-host Environment	3-11
4.	System Software Installation and Configuration	4-1
4.1	Installation	4-2
4.1.1	New Installation	4-2
4.1.2	All Upgrades to RAID Manager 6.22 or 6.22.1	4-2
4.2	Solaris Kernel Driver	4-2
4.2.1	sd_max_throttle Settings	4-3
4.2.2	Generating Additional Debug Information	4-3
4.3	format and lad	4-4
4.3.1	Volume Labeling	4-5

- 4.4 Ghost LUNs and Ghost Devices 4-5
 - 4.4.1 Removing Ghost Drives 4-8
- 4.5 Device Tree Rearranged 4-9
 - 4.5.1 Dynamic Reconfiguration Related Problems 4-10
 - 4.5.1.1 Workaround 4-10
- 4.6 SNMP 4-11
- 4.7 Interaction With Other Volume Managers 4-12
 - 4.7.1 VERITAS 4-12
 - 4.7.1.1 VERITAS Enabling and Disabling DMP 4-12
 - 4.7.1.2 HA Configuration Using VERITAS 4-13
 - 4.7.1.3 Adding or Moving Arrays Under VERITAS 4-13
 - 4.7.2 Solstice DiskSuite (SDS) 4-13
 - 4.7.3 Sun Cluster 4-13
 - 4.7.4 High Availability (HA) 4-13
 - 4.7.5 Quorum Device 4-14

5. Maintenance and Service 5-1

- 5.1 Verifying FRU Functionality 5-2
 - 5.1.1 Disk Drives 5-3
 - 5.1.2 Disk Tray 5-4
 - 5.1.2.1 RSM Tray 5-5
 - 5.1.2.2 D1000 Tray 5-5
 - 5.1.3 Power Sequencer 5-5
 - 5.1.4 SCSI Cables 5-6
 - 5.1.5 SCSI ID Jumper Settings 5-7
 - 5.1.6 SCSI Termination Power Jumpers 5-7
 - 5.1.7 LED Indicators 5-7
 - 5.1.8 Backplane Assembly 5-7
 - 5.1.9 D1000 FRUs 5-7

5.1.10	Verifying the HBA	5-8
5.1.11	Verifying the Controller Boards and Paths to the A3x00/A3500FC	5-8
5.1.12	Controller Board LEDs	5-9
5.1.13	Ethernet Port	5-10
5.2	FRU Replacement	5-10
5.2.1	HBA	5-10
5.2.2	Interconnect Cables	5-11
5.2.3	Power Cords	5-11
5.2.4	Power Sequencer	5-11
5.2.5	Hub	5-12
5.2.6	Controller Card Guidelines	5-12
5.2.7	Amount of Cache	5-13
5.2.8	Battery Unit	5-13
5.2.9	Cooling	5-14
5.2.10	Disk Drives	5-14
5.2.11	Disk Tray	5-14
5.2.12	Midplanes	5-15
5.2.13	Reset Configuration and sysWipe	5-16
5.3	Software and Firmware Guidelines	5-16
5.3.1	Firmware, Software, and Patch Information	5-17
5.3.2	RAID Manager 6 Upgrade	5-18
5.3.3	Firmware Upgrade	5-18
6.	Troubleshooting Common Problems	6-1
6.1	Controller Held in Reset, Causes, and How to Recover	6-2
6.1.1	Reason Controllers Should be Failed	6-2
6.1.2	Failing a Controller in Dual/Active Mode	6-3
6.1.3	Replacing a Failed Controller	6-4

6.1.4	Additional ASC/ASCQ Codes	6-5
6.2	LUNs Not Seen	6-6
6.3	Rebuilding a Missing LUN Without Reinitialization	6-7
6.3.1	Setting the VKI_EDIT_OPTIONS	6-7
6.3.2	Resetting the VKI_EDIT_OPTIONS	6-9
6.3.3	Deleting a LUN With the RAID Manager GUI	6-9
6.3.4	Recreating a LUN With the RAID Manager GUI	6-9
6.3.5	Disabling the Debug Options	6-10
6.4	Dynamic Reconfiguration	6-11
6.4.1	Prominent Bugs	6-12
6.4.2	Further Information	6-12
6.5	Controller Failover and LUN Balancing Takes Too Long	6-12
6.6	GUI Hang	6-13
6.7	Drive Spin Up Failure, Drive Related Problems	6-13
6.8	Phantom Controllers Under RAID Manager 6.22	6-14
6.9	Boot Delay (Why Booting Takes So Long)	6-15
6.10	Data Corruption and Known Problems	6-16
6.11	Disconcerting Error Messages	6-17
6.12	Troubleshooting Controller Failures	6-17
A.	Reference	A-1
A.1	Scripts and man Pages	A-2
A.2	Template for Gathering Debug Information for CPRE/PDE	A-3
A.3	RAID Manager Bootability Support for PCI/SBus Systems	A-4
A.4	A3500/A3500FC Electrical Specifications	A-5
A.5	Product Names	A-7

Figures

FIGURE 2-1 SCSI Bus Length Calculation 2-4

FIGURE 2-2 Fibre Channel Connection With Long Wave GBIC Support 2-4

Tables

TABLE 1-1	A3x00/A3500FC Commandments - Thou Shalt	1-2
TABLE 1-2	A3x00/A3500FC Commandments - Thou Shalt Not	1-2
TABLE 1-3	Web Sites	1-4
TABLE 1-4	Terminal Emulation Functionality	1-6
TABLE 1-5	FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family	1-7
TABLE 1-6	FCOs Affecting the Sun StorEdge RSM 2000/A3x00/A3500FC Product Family	1-11
TABLE 2-1	Server Configuration and Maximum Controller Modules Supported	2-12
TABLE 5-1	Controller Module SCSI ID Settings	5-7
TABLE A-1	A1000 Bootability on PCI-Based Hosts	A-4
TABLE A-2	A1000 Bootability on SBus-Based Hosts	A-4
TABLE A-3	A3x00 Bootability on PCI-Based Hosts	A-5
TABLE A-4	A3x00 Bootability on SBus-Based Hosts	A-5
TABLE A-5	Power Consumption Specifications	A-6
TABLE A-6	Product Name Matrix	A-7
TABLE A-7	NVSRAM Product ID	A-8

Preface

The *Sun StorEdge A1000 and A3x00/A3500FC Best Practices Guide* is intended for use by experienced Sun™ engineering personnel (FE, SE, SSE, and CPRE) who have received basic training on the Sun StorEdge™ A1000, A3x00/A3500FC. It is not intended to replace the existing documentation set, but rather to serve as a single point of reference that provides some answers to questions relating to common installation and service tasks. Further, it serves as a roadmap to more detailed information already provided in the current documentation set and on Sun web sites.

Before You Read This Book

To fully use the information in this document, you must have thorough knowledge of the topics discussed in all of the documents listed in “Related Documentation” on page -xvi.

How This Book Is Organized

This manual is organized as follows:

Chapter 1 introduces some of the tools that are available to help troubleshoot the Sun StorEdge A3x00/A3500FC disk array.

Chapter 2 provides some additional information, guidelines, and tips relating to the installation and configuration of hardware.

Chapter 3 provides some additional information, guidelines, and tips relating to the installation and configuration of RAID Manager.

Chapter 4 provides some additional information, guidelines, and tips relating to installation and configuration of system software.

Chapter 5 provides maintenance and service information for verifying FRU functionality, guidelines for replacing FRUs, and tips on upgrading to the latest software and firmware levels.

Chapter 6 discusses some common problems encountered in the field and provides additional information and tips for troubleshooting.

Appendix A contains the following reference information: a link to the RAID Manager 6.22 README file, a supplementary listing of available man pages for RAID Manager 6.22 commands, a template for gathering debug information, a bootability support matrix, Sun StorEdge A3500/A3500FC electrical specifications, and a product name matrix.

Using UNIX Commands

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

See one or more of the following for this information:

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

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

Shell Prompts

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

Related Documentation

Application	Title	Part Number
Installation and Service	<i>Sun StorEdge A3500/A3500FC Controller Module Guide</i>	805-4980
Installation and Service	<i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i>	805-4981
Installation	<i>Sun StorEdge A3500/A3500FC Task Map</i>	805-4982
Installation and Service	<i>Sun StorEdge A3x00 Controller FRU Replacement Guide</i>	805-7854
Installation	<i>Sun StorEdge A3500FC Controller Upgrade Guide</i>	806-0479
Installation and Service	<i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i>	805-3067
Installation	<i>Sun StorEdge RAID Manager 6.1.1 Update 2 Release Notes</i>	805-3656
Installation and Service	<i>Sun StorEdge RAID Manager 6.1.1 Installation and Support Guide for Solaris</i>	805-4058
Installation and Service	<i>Sun StorEdge RAID Manager 6.1.1 User's Guide</i>	805-4057
Installation and Service	<i>Sun StorEdge RAID Manager 6.22 Installation and Support Guide for Solaris</i>	805-7756
Installation	<i>Sun StorEdge RAID Manager 6.22 Release Notes</i>	805-7758
Installation and Service	<i>Sun StorEdge RAID Manager 6.22 User's Guide</i>	806-0478
Release Notes	<i>Sun StorEdge RAID Manager 6.22.1 Release Notes</i>	805-7758
Installation	<i>Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide</i>	806-7792

Accessing Sun Documentation

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

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

Sun Welcomes Your Comments

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

docfeedback@sun.com

Please include the part number (806-6419-13) of your document in the subject line of your email.

Troubleshooting Overview

This chapter introduces some of the tools that are available to help troubleshoot the Sun StorEdge A3x00/A3500FC disk array, tips for filing a bug, and a listing of the latest field information notices (FINs) and field change orders (FCOs).

This chapter contains the following topics:

- Section 1.1, “A3x00/A3500FC Commandments” on page 1-2
- Section 1.2, “Available Tools and Information” on page 1-3
- Section 1.3, “Tips for Filing a Bug” on page 1-6
- Section 1.4, “FINs and FCOs” on page 1-6

1.1 A3x00/A3500FC Commandments

Tables 1-1 and 1-2 contain PDE recommendations and tips that should be read and followed prior to performing any installation or service tasks on the Sun StorEdge A3x00/A3500FC disk array.

TABLE 1-1 A3x00/A3500FC Commandments - Thou Shalt

Number	Commandment
1	Read the RAID Manager 6 Release Notes and Early Notifier 20029.
2	Only upgrade RAID Manager 6 software and firmware <i>if and only if</i> the controller module, LUNs, and disk drives are all in an optimal state.
3	Never replace the D1000/ESM card while the D1000 is powered on. Refer to the procedure in FIN I0670-1 for proper replacement.
4	Be trained before attempting to use the serial port access.
5	Follow the <i>Sun StorEdge A3500/A3500FC Hardware Configuration Guide</i> for configuration changes between 1x5, 2x7, or 3x15.
6	Always keep the interface hardware, software, firmware, and patches current (refer to Early Notifier 20029).
7	Always reset the battery date on both controllers with <code>raidutil</code> after battery replacement.
8	Always upgrade the controller firmware after a RAID Manager 6 upgrade or installation of RAID Manager 6 patches.
9	Always sync up the controller firmware after a controller board replacement.
10	Follow procedures when replacing disk drives: fail and then revive. See Section 2.3.2, "Adding or Moving Disk Drives to Existing Arrays" on page 2-7.

TABLE 1-2 A3x00/A3500FC Commandments - Thou Shalt Not

Number	Commandment
1	Do not mix RSM2000 (with RSM trays) and A3000 (with D1000 trays) controllers in the same module.
2	Do not downgrade firmware unless the controller is at universal FRU level 2.5.6.32. See FIN I0553 for details.
3	Do not hot swap a controller that "owns" LUNs.
4	Do not move SIMMs from a failed controller to a new controller.

TABLE 1-2 A3x00/A3500FC Commandments - Thou Shalt Not (*Continued*)

Number	Commandment
5	Do not perform <code>boot -r</code> while a controller is held in reset. See Section 6.1, “Controller Held in Reset, Causes, and How to Recover” on page 6-2.
6	Do not enable 16/32 LUN support unless it is necessary (refer to FIN I0589).
7	Do not run A3x00s in a production environment without a LUN 0.
8	Do not move disk drives between hardware arrays (A1000, RSM2000, A3x00, and A3500FC) or in the same array.
9	Do not enable DMP on VxVM pre-3.0.2 releases; RAID Manager 6.22x/DMP compatibility issue has been resolved in VxVM 3.0.2. Refer to the VxVM documentation for details.
10	Do not revive a disk drive if it has been failed by a controller.

1.2 Available Tools and Information

This section contains the following topics:

- Section 1.2.1, “Documentation” on page 1-3
- Section 1.2.2, “Web Sites” on page 1-4
- Section 1.2.3, “Internal Directory” on page 1-4
- Section 1.2.4, “Obtaining the Latest Version of RAID Manager” on page 1-4
- Section 1.2.5, “RAID Manager 6.0, 6.1 and 6.22 Are not Supported” on page 1-5
- Section 1.2.6, “Serial Cable” on page 1-5
- Section 1.2.7, “RAID Manager 6.xx Architecture White Paper Available” on page 1-6

1.2.1 Documentation

The current documentation set is listed in “Related Documentation” on page -xvi.

The documentation set is also available on the following web sites:

`http://edist.central`

`http://infoserver.central/data/syshbk`

1.2.2 Web Sites

The internal and external web sites listed in TABLE 1-3 provide quick access to a wide variety of relevant information.

TABLE 1-3 Web Sites

Web Site Name	URL
Sonoma Engineering	http://webhome.sfbay/A3x00
Network Storage	http://webhome.sfbay/networkstorage
NSTE (QAST) Group	http://webhome.sfbay/qast
OneStop Sun Storage Products	http://onestop.Eng/storage
Enterprise Services Storage ACES	http://trainme.east
Escalation Web Interface	http://sdn.sfbay/cgi-bin/access2
CPRE Group Europe	http://cte-www.uk
Disk Drive Models / FW	http://diskworks.ebay

Note – The Enterprise Services Storage ACES web page requires a login/password for access to certain areas. Information is provided on the home page on how to obtain a login account.

1.2.3 Internal Directory

The following internal directory contains the released versions of RAID Manager 6 software, firmware and NVSRAMs:

```
net/artemas.ebay/global/archiva/archive/StorEdge_Products  
/sonoma
```

1.2.4 Obtaining the Latest Version of RAID Manager

RAID Manager 6.22x management software for the A1000, A3x00, and A3500FC disk arrays is available for download via the Sun Download Center (formerly called the Sun Software Shop). It is under the link Download at:

```
http://www.sun.com/storage/disk-drives/raid.html
```

Note – Before you attempt to install the RAID Manager software, be sure to read the *Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide* and Early Notifier 20029 for the latest installation and operation procedures.

1.2.5 RAID Manager 6.0, 6.1 and 6.22 Are not Supported

RAID Manager 6.0 and 6.1 have been superseded by newer versions of RAID Manager. RAID Manager 6.1.1 is only supported in cases involving data corruption or loss. Upgrade to RAID Manager 6.22.1 as soon as possible.

1.2.6 Serial Cable

Sun field service personnel who have undergone the proper training on serial port access can obtain a serial port cable along with the *Debug Guide* from the Area Technical Service Manager. If you need assistance with serial port functions, contact the local Storage ACES or CPRE. CPRE has access to LSI Logic's 24 hour support line.

The serial port provides access to useful commands used to determine controller, drive, and LUN status. It was originally intended to be used by developers.



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

The serial cable is an LSI Logic, Inc. proprietary DB-15 to DB-25 cable. The serial cable comes with a DB-25 to DB-9 extension for PC interface. To connect to the serial port on an UltraSPARC™ machine you will need a DB-25 gender adapter. The gender adapter is not required for pre-UltraSPARC machines.

Note – Do not leave the serial port cable and *Debug Guide* with the customer. This cable is for use by trained Sun personnel only.

If you use a PC to connect to the serial port of the disk array, you need terminal emulation software. Also, you need to ensure that the `Break` functionality is available. Although there are many different software applications that provide terminal emulation, you will have the best results if you use the applications listed in TABLE 1-4.

TABLE 1-4 Terminal Emulation Functionality

Operating Environment	Communication Software
Microsoft Windows 98/2000	Procomm Plus
Solaris x86/SPARC™	tip
Linux	dip

1.2.7 RAID Manager 6.xx Architecture White Paper Available

The white paper titled *StorEDGE A3000/A1000 Controller Architecture (The BlackBox behind the RAID Module)* is available on the Sonoma Engineering web site:

http://webhome.sfbay/A3x00/HW/A3000_controller_paper.fm.ps

1.3 Tips for Filing a Bug

Refer to Section A.2, “Template for Gathering Debug Information for CPRE/PDE” on page A-3 for a template that should be used when filing a RAID Manager/A3x00FC related bug.

1.4 FINs and FCOs

For access to the latest field information notices (FINs) and field change orders (FCOs), refer to the following web site:

http://sdpsweb.ebay/FIN_FCO

Tables 1-5 and 1-6 list the current FINs and FCOs affecting the Sun StorEdge RSM 2000/A3x00/A3500FC product family as of January 2001.

TABLE 1-5 FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family

FIN Number	Release Date	Product	Description
I0310	07/09/97	SSA RSM 21x RSM Array 2000	Updated - Failure to follow documented installation procedures to remove shipping brackets and reseal drives may cause multiple disk errors.
I0312	08/31/00	RSM Array 2000 patches	Required set of patches for the Sun RAID Storage manager Array 2000.
I0324	05/06/97	Sun VTS, RSM Array 2000	Problem with Sun VTS not probing LUNs located on an RSM Array 2000 system.
I0325	06/09/97	RSM Array 2000, Power sequencer	Power in the RSM Array 2000 will be on when the keyswitch is off and one of the power sequencer's circuit breaker is off.
I0328	06/13/97	RSM Array 2000 disk configuration	Some RSM 2000's installed with 9.0 GB disk drives were configured as if they only had 4.0 GB disk drives attached.
I0332	06/20/97	SSA 21x, RSM Array 2000	INFORMATIONAL ONLY! Disk LEDs light or blink in an inconsistent manner during power up on some disk trays.
I0334	07/09/97	SSA RSM Model 219, RSM Array 2000 Seagate ST19171WC 9.0 GB disk drive	"Cushioning damper" kit for Seagate ST19171WC 9.0 GB disk drives installed in a SPARCstorage Array RAID Storage Manager Model.
I0362	12/18/98	RSM Array 2000 RM 6.1	UPDATED FIN. Updated parity error recovery procedure.
I0368	04/10/98	RSM Array 2000 controller board	UPDATED FIN. Memory error detection on the RSM Array 2000 controller logic board is ineffective.
I0440	10/29/98	RM Rev 6.1.1	RM 6.1.1 with Sun StorEdge A1000, A3000, or A3500 may experience status check problems.
I0441	10/30/98	Sun StorEdge D1000 disk tray	New disk firmware revision 7063 for ISP boot failures on D1000's with Seagate ST39173WC 9 Gbyte disks.
I0473	03/09/99	Solaris 7, RM6.1.1 UG A3x00/A1000	Solaris 7 upgrade recommendations for Sun StorEdge A3x00/A1000.
I0505	08/09/99	RM 6.1.1 RAID LUN recovery	UPDATED FIN. A3x00/A1000 RAID 0 LUN recovery requires stopping I/O to the LUN.
I0509	06/30/99	RM 6 and Solaris kernel corruption	RAID Manager large rdriver.conf files need a patch to prevent panic.

TABLE 1-5 FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family

FIN Number	Release Date	Product	Description
I0511	07/28/99	VM DMP interferes with A3x00 RDAC	Enabling DMP with RDAC on A3x00 and A1000 may cause private regions to be lost.
I0520	04/24/01	Quorum device in Sun Cluster	Servicing storage that contains a device that is used as a quorum device in a Sun Cluster environment.
I0531	01/07/00	A3500FC parallel LUN modifications	UPDATED FIN. Sun StorEdge A3500FC with preconfigured LUNs may encounter errors if parallel LUN modifications are made using Sbus host adapters.
I0536	09/09/00	E10000 with A3x00 DR	E10000 systems with an A3x00 attached may encounter DR errors.
I0594	09/08/00	Sun StorEdge A3000 name plate	A3000 name tag was applied on the product instead of Sun StorEdge A3500 causes confusion to the customers.
I0547	01/21/00	UDWIS fcode	Intermittently, SCSI devices connected to a UDWIS card may not be usable after a reboot.
I0551	02/08/00	Large Sun StorEdge A1000, A3000, A3500 configurations	Boot process and controller on-line process may take hours in systems with large Sun StorEdge A1000, A3000, A3500, or configurations.
I0552	01/24/00	isp driver bug	SCSI devices (especially in a multi-hosted configuration) may go off-line after isp errors.
I0553	02/04/00	A3x00 controller firmware	The firmware download procedure documented in A3x00 Controller Replacement Guide may render the controller unusable.
I0557	05/23/00	RM 6.1.1 firmware	Differences in LUN capacity after firmware upgrade on Sun StorEdge A3000 or RSM 2000 hardware RAID controllers.
I0566	06/12/00	Solaris 8 with any Sun StorEdge A3x00/A1000	A patchID 108553-xx is required to run Solaris 8 with RM 6.22.
I0569	03/31/00	isp fcode 1.28 probe error	UDWIS Sbus host adapter with isp-fcode 1.28 can cause "invalid command" during probe-scsi-all or boot -hv from a Sun StorEdge D1000.
I0573	09/28/00	Storage Array A1000, A3x00	Sun StorEdge A1000, A3000, A3500, or A3500FC requires the existence of LUN 0 for proper operation.
I0579	06/07/00	UDWIS/Sbus SCSI adapter	Systems with Ultra DWIS/Sbus host adapter slow to a halt under heavy I/O loads and the console displays "SCSI transport failed" error messages.
I0586	06/22/00	Sun StorEdge A3500FC connectivity	The hardware and software requirement to use the onboard SOC+ the I/O boards to connect the A3500FC.

TABLE 1-5 FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family

FIN Number	Release Date	Product	Description
I0589	06/21/00	Any RM 6 version	glm.conf file must be modified to support more than 8 LUNs on any PCI-SCSI connected A1000 or A3x00 using any version of RM 6.
I0590	07/20/00	Sun StorEdge A3500FC upgrade kit	The SCSI NVSRAM will overwrite the FC controller NVSRAM if following the documented A3500 SCSI to FC upgrade procedure.
I0594	I0594	Sun StorEdge A3x00 and A3500 arrays.	Some StorEdge A3500 products shipped with the StorEdge A3000 name tag lead customers to believe the wrong product was delivered.
I0612	09/08/00	Sun StorEdge Array Configuration	RM6 may detect a non-existent 2MB dacstore on factory formatted hard disk drives.
I0613	09/20/00	Sun StorEdge A1000 array	When replacing a failed controller board on the A1000, it is important to first verify the version of RAID Manager software, the controller firmware level and controller memory size.
I0619	10/12/00	Sun StorEdge A1000 and A3x00 Arrays	Proper procedures for booting from the Sun StorEdge A1000 or A3x00 hardware RAID devices, including known issues and problems.
I0634	11/30/00	Sun StorEdge A3x00 Array controller failover	Sun StorEdge A3x00 Array controller failover may cause temporary delay in processing of pending disk I/O.
I0637	01/10/01	Sun StorEdge A3500 Array	A3500 controller powering up prior to Dilbert trays due to wrong connections of L5, R5,L6, and R6 power sequencer cables.
I0643	01/10/01	Sun StorEdge A3x00 with RAID Manager	Sun StorEdge A3x00 array configurations with RAID Manager 6.1.x may be susceptible to controller “deadlocks”.
I0648	02/16/01	Sun StorEdge A3x00 Array	Disks in Sun StorEdge A3x00 arrays might go offline, making the devices unavailable.
I0653	05/08/01	RSM 2000 array shelves	RSM array shelves may collapse under weights greater than 40 lbs. (18.2 kg), causing system damage or personal injury.
I0670	06/15/01	Sun StorEdge D1000 ESM board and Sun StorEdge A3x00/A3500FC arrays	Procedure to replace the D1000 Environmental Service Module (ESM) board in A3x00/A3500FC arrays.
I0684	06/21/01	RAID Manager 6.22 healthchk utility	RAID Manager 6.22 healthchk utility might not report power supply or fan failures in Sun StorEdge A3x00 arrays, which might result in loss of availability.

TABLE 1-5 FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family

FIN Number	Release Date	Product	Description
I0685	07/18/02	RAID Manager 6 Software	Certain precautions need to be observed when upgrading the Solaris operating environment on systems with RAID Manager 6 software installed.
I0688	06/27/01	RAID Manager 6.22 on Solaris 2.5.1	NSTE (QAST) qualified RAID Manager 6.22 to run on Solaris 2.5.1.
I0698	07/12/01	RAID Manager 6.22 patch ID 108834-09 and 108553-09	Installing RAID Manager 6.22 patches 108834-09 and 108553-09 might generate false or misleading warning messages "unresponsive drives" when running the healthchk or drivutil commands and report excessive false 9501 messages.
I0709	03/19/02	Sun StorEdge A1000/A3x00/A3500FC arrays	After replacing a controller FRU on A1000 / A3x00/A3500FC with RM 6.22.1, the NVSRAM must be reloaded.
I0724	10/10/01	Sun StorEdge A1000/D1000 and A3500/A3500FC arrays	18.2 Gbyte and 36 Gbyte IBM disk drives may be susceptible to early life failures.
I0727	10/17/01	Sun StorEdge A1000/A3x00 controllers	Recovering A1000 / A3x00 controller C numbers after a device path changed due to reboot -r.
I0736	10/26/01	Sun StorEdge A3500FC controller	Current replacement procedures for an A3500FC controller in a clustered environment could result in the controller going off line.
I0738	04/17/02	Sun StorEdge A3x00 array with RAID Manager	Access to raw partitions on Sun StorEdge A1000, A3x00, or A3500FC LUNs by non-root users, such as Oracle or Sybase, is not allowed with some patch levels of RAID Manager 6.22 and 6.22.1.
I0744	11/30/01	RAID Manager 6.22.1	RM 6.22.1 problems reported in BugId 4521759 have been fixed with Patches 112126-01 for Solaris 8 and 112125-01 for other Solaris versions above Solaris 2.5.1.
I0782	03/01/02	Sun StorEdge A1000/A3x00 arrays	Adding new disk drives to StorEdge A1000 or A3X00 could cause RAID Manager to become unavailable.
I0786	03/04/02	Sun StorEdge A1000/A3x00 arrays	'hot_add' of an A1000/A3x00/A3500FC array may need to be followed by a reconfiguration reboot on Solaris 8.
I0809	04/10/02	RM 6.22.1	A Solaris 2.6 system with /usr on separate slice may fail to boot after upgrade to RM 6.22.1.
I0825	05/10/02	RM 6.22.1 on A1000/A3x00/A3500FC	Running automatic parityck on A1000/A3x00/A3500FC arrays with RAID Manager 6.22.1 does not repair problems.

TABLE 1-5 FINs Affecting the Sun StorEdge A1000 and RSM 2000/A3x00/A3500FC Product Family

FIN Number	Release Date	Product	Description
I0828	05/21/02	RM 6.22.1	LUNs may become inaccessible after upgrading from RM 6.22 to 6.22.1 or after adding unformatted disk drives to RM 6.22.1.
I0845	06/25/02	RAID Manager 6	RAID Manager 6 may hang for 3-8 minutes when an IBM drive is in the failed state in a Sun StorEdge A1000/A3x00/A3500FC array.
I0860	08/02/02	RAID Manager 6 with Solaris 9	RAID Manager 6.22.1 with A1000/A3x00/A3500FC arrays needs a new patch revision for Solaris 9.
I0865	TBD	Sun StorEdge A3500FC Arrays	Extraneous physical device paths may appear.

TABLE 1-6 FCOs Affecting the Sun StorEdge RSM 2000/A3x00/A3500FC Product Family

FCO Number	Release Date	Product	Description
A0120	02/27/98	RSM Array 2000 WD2S Card	Power glitch may cause disks to go off-line on RSM Array 2000 configured without an uninterruptable power supply (UPS).
A0162	03/09/00	Power Supply A1000/D1000	Sun StorEdge A1000 and D1000 power supplies fail with amber LED.
A0163	03/31/00	UDWIS Sbus host adapter A1000/D1000, A3x00/A3500	Systems with large quantities of UDWIS/Sbus host adapters installed may not come up after reboot due to miscommunication between the SCSI host and the target.
A0164	05/12/00	FC100/P card, A5x00/A3500	Systems with FC100/P cards (375-0040-xx) having Molex optical transceivers, may cause increasing numbers of loop errors, excessive LIPs, and SCSI parity errors as they age.
A0165	06/19/00	RSM 2000 A3000 battery backup unit	Any battery backup unit over two years old may not have enough energy to hold data in the RSM 2000 or A3000 controller's write-back cache for 72 hours in the event of a power outage.

Hardware Installation and Configuration

This chapter provides some additional information, guidelines, and tips relating to the installation and configuration of hardware.

This chapter contains the following sections:

- Section 2.1, “New Installation” on page 2-2
- Section 2.2, “Adding or Moving Arrays to a Host With Existing A3x00 Arrays” on page 2-6
- Section 2.3, “Adding Disks or Disk Trays” on page 2-6
- Section 2.4, “Setting Up 2x7 and 3x15 Configurations, and Converting 1x5 to 2x7 or 3x15” on page 2-8
- Section 2.5, “Sun StorEdge A3500/A3500FC Lite” on page 2-9
- Section 2.6, “Cluster, Multi-Initiator, and SCSI Daisy Chaining Configurations” on page 2-10
- Section 2.7, “Supported Configurations” on page 2-11
- Section 2.8, “SCSI to FC-AL Upgrade” on page 2-14

2.1 New Installation

This section contains the following topics:

- Section 2.1.1, “Battery Unit” on page 2-2
- Section 2.1.2, “Power Cables” on page 2-2
- Section 2.1.3, “Power Sequencer” on page 2-3
- Section 2.1.4, “Local/Remote Switch” on page 2-3
- Section 2.1.5, “SCSI and Fiber-Optic Cables” on page 2-3
- Section 2.1.6, “SCSI ID, Loop ID, Controller, and Disk Tray Switch Settings” on page 2-4
- Section 2.1.7, “World Wide Name (WWN)” on page 2-5

2.1.1 Battery Unit

For further information regarding the battery unit, see Section 1.3.1 “Battery Unit” in the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Verify that the battery’s date of manufacture is not older than six months. The battery’s date of manufacture can be found on the Battery Support Information label located on top of the battery. Write down the date of installation on the Battery Support Information label. Add two years to the date of installation and write down the replacement date on the Battery Support Information label.

The battery unit will undergo self diagnostics at power up which will take approximately 15 minutes to complete. During this time you might see the following SCSI message: `ASC/ASCQ A0/00` which indicates that write cache cannot be enabled. This message should go away within 15 minutes or so unless the battery unit was completely drained. In this case it will require additional time for the battery to completely recharge.

2.1.2 Power Cables

For further information and guidelines for connecting the AC power cables, refer to Chapter 3 in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*. It is important that you follow these guidelines.

The controller must to be powered on at the same time or after the disk trays so that the controller registers the disk drives. The power cables to the controller should always be connected to one of the four outlets on the second sequenced group of the

expansion cabinet power sequencer. If the cabinet's original factory configuration has not been changed, then the cabinet should contain the correct power sequencer connections.

Note – At the bottom of the expansion cabinet are two power sequencers. The front power sequencer is hidden behind the front key switch panel. Remove the front key switch panel to gain access to the power sequencer's power cable.

2.1.3 Power Sequencer

For further information regarding power sequencer configuration, refer to Chapter 3 in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.



Caution – To prevent the possibility of data loss due to an inadvertent power shutdown of an expansion cabinet, ensure that in a 3x15 configuration the power sequencers are daisy chained front to front and back to back. The Local/Remote switch should be set to Remote.

2.1.4 Local/Remote Switch

The Local/Remote switch on each power sequencer is factory set to Remote (default). This allows power on/off control of each power sequencer through the front key switch. If the Local/Remote switch is set to Local then the power on/off control of each power sequencer is controlled by each power sequencer's main power circuit breaker switch.

2.1.5 SCSI and Fiber-Optic Cables

For further details regarding the SCSI and/or fiber-optic cables, refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

Each A3x00 is shipped from the factory with two 12-meter SCSI host cables. The recommended maximum SCSI bus length is 25 meters. Exceeding this length can cause SCSI or data errors. When calculating the maximum SCSI bus length in a particular configuration, remember to include the internal SCSI bus of each device which is 0.1 meters.

The example configuration in FIGURE 2-1 shows a daisy chained multi-initiator configuration consisting of two hosts connected to two A3x00 modules.

The total SCSI bus length for this example is 24.4 M. It is calculated as follows: each SCSI cable's length (Cable no. 1 + Cable no. 2 + Cable no. 3) + the internal SCSI bus length of each device (Host no. 1 + A3x00 no. 1 + A3x00 no. 2 + Host no. 2).




FIGURE 2-1 SCSI Bus Length Calculation

Each A3500FC is shipped from the factory with one pair of 15 meter fiber-optic cables. The product will support a multi-mode fiber-optic connection up to 500 meters long using short wave GBICs only. The use of long wave GBICs is supported but the connection must be through a switch to the host processor (FIGURE 2-2). Refer to the switch documentation for the maximum cable length for the long wave GBIC connection.

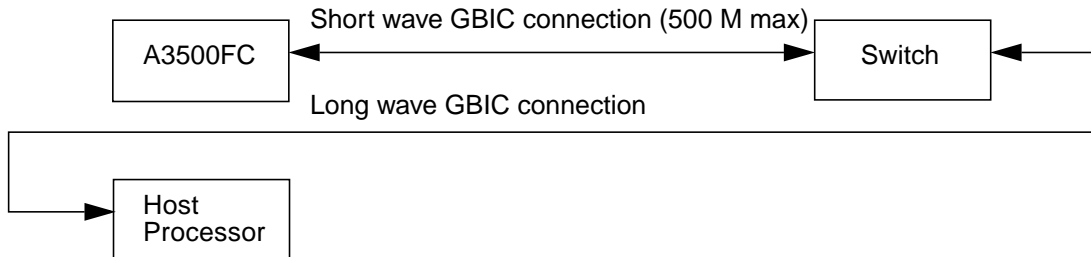


FIGURE 2-2 Fibre Channel Connection With Long Wave GBIC Support

2.1.6 SCSI ID, Loop ID, Controller, and Disk Tray Switch Settings

The factory default setting for controller A is SCSI ID 5 (T5) and for controller B is SCSI ID 4 (T4). The SCSI cables connected within a SCSI daisy chain configuration should be crossed to avoid SCSI ID conflicts (see Section 1.2.2 in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*).

If two disk trays have the same tray ID, the system reports a 98/01 ASC/ASCQ error code during boot up time. In the 1x2 configuration since two drive channels from a controller share one disk tray, it is unavoidable to have a tray ID conflict. During boot up time, the 98/01 ASC/ASCQ error code is reported but has no impact on system performance.

In a Fibre Channel dual controller connection through two hubs, each hub should be connected to one “A” controller and one “B” controller to avoid SCSI ID conflicts (see Section 2.2.4 in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*).

Also, see Section 2.3 “Setting the Loop ID” in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide* for procedures to set the loop ID for A3500FC controllers. This section also contains a detailed SCSI ID/loop ID reference chart.

For better redundancy, do not connect both controllers of the same controller module to the same I/O board.

It is recommended that you set both controllers to active/active mode and to balance the LUNs across both controllers for better performance. The maximum number of supported LUNs is the total number of LUNs between the two controllers. For example, if the maximum supported number of LUNs is 16 then you can have 8 LUNs on controller A and 8 LUNs on controller B. If one controller becomes off-line then the surviving controller owns all 16 LUNs.

2.1.7 World Wide Name (WWN)

Each FC-AL device requires a unique World Wide Name (WWN) that identifies the device or controller on the loop. There is no hardware setting for the WWN. The WWN is a 16-byte hexadecimal value that can be retrieved as the logical unit WWN in the device identifier inquiry page (0x83) and is generated as follows:

- The name (derived from the controller MAC address) of the storage array controller creating the volume in the upper 8-bytes.
- An auto increasing number and time stamp in the lower 8-bytes.

Note – When a controller is hot swapped, the upper 8-bytes of the WWN may not match a controller in the storage array.

2.2 Adding or Moving Arrays to a Host With Existing A3x00 Arrays

When moving a disk array, ensure that the array being moved has firmware levels that match with the new host. See "Upgrading Controller Firmware" in the *Sun StorEdge RAID Manager 6.22 Release Notes*. Since the firmware on the controller cannot be downgraded, except in the case of a universal FRU, you should not move an array to a host with a lower RAID Manager release. If the array being moved has more than 8 LUNs, be sure the new host has greater than 8 LUN support turned on. See "Maximum LUN Support in Solaris 2.6 and Solaris 7 Environments" in the *Sun StorEdge RAID Manager 6.22 Release Notes*. Also, refer to "Adding or Moving Arrays Under VERITAS" on page 4-13 for information pertaining to VERITAS volumes.

You can move an A1000 or an A3x00 array from one server to another and still be able to access the data after the move.

You can perform most but not all of the RAID Manager 6 array management commands after the move. In particular, you will not be able to execute commands such as LUN creation or deletion. This can only be done on the host where the array was originally connected. This is because there is host ownership information stored in DacStore and on the LUNs that prevent these types of commands from being executed. The workaround involves using either the storutil command to update the host ownership information in the DacStore region or the RAID Manager 6 GUI Select Module -> Edit function to change the module name both of which updates the DacStore with the new owning host information. See the man page storutil (1M) to view the available command line options.

2.3 Adding Disks or Disk Trays

This section contains the following topics:

- Section 2.3.1, "Adding or Moving Disk Trays to Existing Arrays" on page 2-7
- Section 2.3.2, "Adding or Moving Disk Drives to Existing Arrays" on page 2-7

2.3.1 Adding or Moving Disk Trays to Existing Arrays

RAID Manager 6.22 has dynamic capacity expansion capability. If your RAID system has not used all five drive side channels, you can add disk trays to it and expand the capacity of the existing drive groups. The existing LUN capacity does not increase. See “Configuring RAID Modules” in the *Sun StorEdge RAID Manager 6.22 User’s Guide*. When expanding to a different configuration, ensure that it is a supported configuration as documented in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

Note – Disk drives should not be interchanged between A1000, A3500, and A3500FC. The controller may update its NVSRAM from the (inconsistent) DacStore in the foreign disks causing an A3x00/A3500FC controller to start behaving like an A1000 or RSM 2000 controller. If this should occur you will need to reload the correct NVSRAM into each controller.

The A3000 back-end SCSI transfer rate is set to 20-MB/sec and the A3500 back-end SCSI transfer rate is set to 40-MB/sec. Interchanging of these controllers can cause data corruption problems.

You need to power cycle the controller module for the new disk tray to be recognized in. Be sure the drives contained in the new disk tray are not larger in capacity than the Global Hot Spare (GHS) drives. This is because the GHS will not spare a drive that has a higher capacity than itself. To connect the power and SCSI cables to the new disk tray, and to set the disk tray ID (and the option switch on D1000 trays), refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

2.3.2 Adding or Moving Disk Drives to Existing Arrays

You can add new drives to empty slots in your existing array. To add new disk drives:

1. **Insert a disk drive into the tray.**
2. **Allow time for the drive to spin up (approximately 30 seconds).**
3. **Run a health check to insure that the drive is not defective.**
4. **Fail the drive then revive the drive to update DacStore on the drive.**
5. **Repeat steps 1 through 4 until all drives have been added.**
6. **New drives are now ready for LUN creation.**

Note – Refer to Escalation no. 525788, bug 4334761. Refer to FIN I0612 for further information.



Caution – If the drives you are adding to the array were previously owned by another controller module, either A1000 or A3x00/A3500FC, ensure that you pre-format the disk drives to wipe clean the old DacStore information before inserting them in an A3x00/A3500FC disk tray.



Caution – Do not randomly swap drives between drive slots or RAID systems. You must use the Recovery Guru procedure to replace drives. Also see Section 5.2.10, “Disk Drives” on page 5-14.



Caution – If you take out a disk drive, put a dummy drive in so proper cooling can be maintained.

2.4 Setting Up 2x7 and 3x15 Configurations, and Converting 1x5 to 2x7 or 3x15

To connect the power and SCSI cables to the new disk tray, and to set the disk tray ID (and the option switch on D1000 trays), refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

Because of the re-cabling, the drive DacStore will be erroneous. Some drives might fail or not appear. Ghost drives might appear. Do not attempt to fix these problems. You first need to delete all previous LUN configurations. A Reset Configuration may work, but the best thing to do is issue a `sysWipe` command then `sysReboot` on each controller through the serial port.



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

Note – When you issue a `sysWipe` command you might see a message indicating that `sysWipe` is being done in a background process. Wait for a message indicating that `sysWipe` is complete before issuing a `sysReboot` command. Once the configuration is reset and the previous `DacStore` is cleaned up, the drive status should come up Optimal as long as the drive has no internal problem. `sysWipe` should be run from each controller.

Note – An A3000 RAID system cannot be converted to an A3500 RAID system by swapping disk trays because the controller NVSRAM is different. The A3x00 requires an entire controller module upgrade. Also, D1000 trays are not supported in the A3000 56" cabinet. You will need to upgrade the cabinet.

2.5 Sun StorEdge A3500/A3500FC Lite

This package includes two rackmountable Sun StorEdge D1000s, a RAID module, and parts necessary to connect to a host server.

Refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide* for cabling instructions. Be sure the RAID controller is connected to power cords from sequenced switch group 2.

Refer to the following web site for ordering information:

<http://store.sun.com/docs/specials/bundlegroup.jhtml?catid=48293>

Note – Since two drive channels share a disk tray in this configuration, it is recommended for a RAID 1 LUN configuration to mirror between two drive trays. If you lose one D1000 disk tray, you will lose two drive channels.

The A3500/A3500FC Lite disk trays were previously available with 9-GB disk drives. Starting September 2000 only 18-GB disk drives will be available for purchase.

2.6 Cluster, Multi-Initiator, and SCSI Daisy Chaining Configurations

This section contains the following topics:

- Section 2.6.1, “Cluster Information” on page 2-10
- Section 2.6.2, “Multi-Initiator Information” on page 2-11
- Section 2.6.3, “SCSI Daisy Chaining Information” on page 2-11

2.6.1 Cluster Information

Refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide* for instructions about cabling and setting up the SCSI ID and/or the FC loop ID.

The Sun Cluster home page is located on the following web site:

`http://suncluster.eng/index.html`

For detailed cluster configuration information, refer to the *Sun Enterprise Cluster System Hardware Site Preparation, Planning, and Installation Guide*. You can find this document on the following web site:

`http://suncluster.eng.sun.com/engineering`

For more details about controller replacement and restoring in a cluster environment, refer to the *Sun Cluster 3.0 Hardware Guide*, “How to Replace a Failed StorEdge A3500 Controller or Restore an Offline StorEdge A3500 Controller”. Sun Cluster 3.0 documents are available on the following web site:

`http://suncluster.eng/products/SC3.0`

Specific Sun Cluster documents and relevant chapters within these documents are:

- Sun Cluster 2.2—Sun Enterprise Cluster 3000/4000/5000/6000/10000 Hardware Service Manual, 805-6512, chapter 3 “Hardware Troubleshooting and chapter 9 “Major Subassembly Replacement”.
- Sun Cluster 3.0—Sun Cluster 3.0 Hardware Guide, 806-1420, chapter 7 “Installing, Configuring, and Maintaining a Sun StorEdge Disk Array”.

2.6.2 Multi-Initiator Information

Hubs are required for connecting A3500FCs in cluster/multi-initiator configurations. A3x00/A3500FC is supported with Sun Cluster 2.2 in a two node cluster configuration.

In a multi-initiator (aka multi-host) connection, both nodes need to be Sun SPARC servers. Only a multi-initiator configuration that runs Sun Cluster is supported by Sun. This applies to A3x00 and A3500FC.

Refer to the following web site for details regarding a cluster support matrix:

<http://suncluster.eng/support-matrix>

Also refer to the following cluster FAQ web site:

<http://suncluster.eng/sales/faq/#storage>

2.6.3 SCSI Daisy Chaining Information

For further information regarding daisy chaining SCSI disk arrays, refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*. Ensure that the SCSI cables connected from controller A to controller B are crossed to avoid SCSI ID conflict.

The equivalent to SCSI daisy chaining within an FC-AL loop configuration is a hub connection. For further information on setting up a hub connection and the controller loop ID, refer to the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*.

2.7 Supported Configurations

This section contains the following topics:

- Section 2.7.1, “Maximum Server Configurations” on page 2-12
- Section 2.7.2, “Onboard SOC+” on page 2-12
- Section 2.7.3, “Second Port on the SOC+ Card” on page 2-13
- Section 2.7.4, “Disk Drive Support Matrices” on page 2-13
- Section 2.7.5, “Independent Controller/Box Sharing” on page 2-13
- Section 2.7.6, “HBAs” on page 2-13

2.7.1 Maximum Server Configurations

Table 2-1 lists the maximum number of controller modules that are supported for a given server configuration.

TABLE 2-1 Server Configuration and Maximum Controller Modules Supported

Server	A3500FC Maximum Controller Modules (Direct)	A3500FC Maximum Controller Modules (Hub)	A3500 SCSI Maximum Controller Modules (Single)	A3500 SCSI Maximum Controller Modules (Daisy- Chained)
Sun Enterprise™ 10000	21	34	21	21
Sun Enterprise 6000/6500	21	34	21	36
Sun Enterprise 5000/5500	10	34	10	20
Sun Enterprise 4000/4500	10	34	10	20
Sun Enterprise 3500	6	24	6	12
Sun Enterprise 3000	4	16	4	8
Sun Enterprise 450	3	12	3	4
Sun Enterprise 250	2	8	2	2
Sun Ultra 80	2	8	Not supported	Not supported
Sun Ultra 60	2	8	Not supported	Not supported

2.7.2 Onboard SOC+

Onboard SOC+ is supported with A3500FC on host platforms: E3x00 through E6x00.

The following types of I/O boards have onboard SOC+ that are supported:

- I/O board with SOC+:
 - X2611 (501-4266-06) I/O type 4, 83 MHz Gigaplane.
 - X2612 (501-4883-05) I/O type 4, 83/90/100 MHz Gigaplane.
- I/O Graphic board with SOC+:

- X2622 (501-4884-05) I/O type 5, 83/90/100 MHz Gigaplane.

Both SOC+ connections on one I/O board can be used to connect to an A3500FC concurrently. For better redundancy, do not connect both controllers of the same controller module to the same I/O board.

Minimum firmware requirement for supported I/O board is: 1.8.25.

Note – Refer to FIN I0586-1 for details.

2.7.3 Second Port on the SOC+ Card

Other Fibre Channel devices can be connected to the second port. Refer to FIN I0586-1 for details.

2.7.4 Disk Drive Support Matrices

The following URL sites provide support information for all hard disk drives that are used in the Sun StorEdge™ RSM Array 2000 and Sun StorEdge A1000, A3x00 and A3500FC disk arrays.

- <http://webhome.sfbay/harddrivecafe/>

2.7.5 Independent Controller/Box Sharing

Independent Controller is also known as *Box Sharing*. It enables the storage capacity of an A3x00/A3500FC disk array to be shared by two independent hosts. See the “Independent Controller Configuration” section in Chapter 2 in the *Sun StorEdge RAID Manager 6.22 User’s Guide* for an overview of this function.

2.7.6 HBAs

Refer to the HBA support matrix in Early Notifier 20029 for a list of supported HBAs.

2.8 SCSI to FC-AL Upgrade

Note – Refer to FIN I0590 and the latest version of the *Sun StorEdge A3500/A3500FC Controller Upgrade Guide* for more information regarding this procedure. The latest version of the *Sun StorEdge A3500/A3500FC Controller Upgrade Guide* at the time this document was prepared: part no. 806-0479-11.

You need to load NVSRAM to the A3500FC controllers as documented in FIN I0590 and in the *Sun StorEdge A3500/A3500FC Controller Upgrade Guide*. Be sure to load the correct NVSRAM patch:

- A3500 array with D1000 trays - patch no. 109232-01
- A3000 array with RSM trays - patch no. 109233-01

If you have RM 6.22.1, then use the NVSRAM in that release, which is later than the NVSRAM in these two patches.

Follow the instructions provided in the *Sun StorEdge A3500/A3500FC Controller Module Guide* to route fiber cables. The old SCSI controller front bezel does not give adequate room for fiber-optic cable routing; discard it. Use the new A3500FC controller face plate that comes with the A3500FC upgrade kit.

RAID Manager Installation and Configuration

This chapter provides some additional information, guidelines, and tips relating to the installation and configuration of an array.

This chapter contains the following sections:

- Section 3.1, “Installation and Configuration Tips, Tunable Parameters, and Settings” on page 3-2
- Section 3.2, “LUN Creation/RAID Level” on page 3-5
- Section 3.3, “LUN Deletion and Modification” on page 3-9
- Section 3.4, “Controller and Other Settings” on page 3-9

3.1 Installation and Configuration Tips, Tunable Parameters, and Settings

This section contains the following topics:

- Section 3.1.1, “Software Installation” on page 3-2
- Section 3.1.2, “Software Configuration” on page 3-3
- Section 3.1.3, “RAID Module Configuration” on page 3-3
- Section 3.1.4, “Tunable Parameters and Settings” on page 3-3
- Section 3.1.5, “Multi-Initiator/Clustering Environment” on page 3-4
- Section 3.1.6, “Maximum LUN Support” on page 3-5

3.1.1 Software Installation

- Upgrading from RAID Manager 6.0/RAID Manager 6.1 to RAID Manager 6.1.1 Update 2 or RAID Manager 6.22x.
- Two Mbyte DacStore is not supported on RAID Manager 6.22x.
Issue: existing LUNs with a 2-MB DacStore running under RAID Manager 6.1.1/RAID Manager 6.22x see FIN I0557 for recommendations and procedures.
- Upgrading from RAID Manager 6.1 to RAID Manager 6.1.1 Update 2.
See the “Sun StorEdge RAID Manager Upgrade Procedure” section in the *Sun StorEdge RAID Manager 6.1.1 Update 2 Release Notes*.
- Installing RAID Manager 6.1.1 Update 2.
Refer to the *Sun StorEdge RAID Manager 6.1.1 Installation and Support Guide for Solaris*.
- Installing RAID Manager 6.22x.
Refer to the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide for Solaris* and the *Sun StorEdge RAID Manager 6.1.1 Update 2 Release Notes*.
Also, refer to the *Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide* (part number 806-7792) for additional information when upgrading to RAID Manager 6.22x.
- Ensure that you also install the latest controller firmware revision after upgrading the RAID Manager software.

3.1.2 Software Configuration

- RAID Manager 6.1.1—Refer to the *Sun StorEdge RAID Manager 6.1.1 Installation and Support Guide* for details.
- RAID Manager 6.22—Refer to the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide* for details.
- If the default LUN 0 has to be resized (remove and recreate because the size is too small), see FIN I0573 for procedure.
- When upgrading to RAID Manager 6.22, you may see warning messages indicating that there are bad disk drives. This is normal and is a new feature in RAID Manager 6.22 called Predictive Failure Analysis (PFA). Refer to Chapter 6 “Recovery” in the *Sun StorEdge RAID Manager 6.22 User’s Guide* for details.

3.1.3 RAID Module Configuration

There are three typical RAID Module configurations (refer to Chapter 2 in the *Sun StorEdge RAID Manager 6.1.1 User’s Guide* or *Sun StorEdge RAID Manager 6.22 User’s Guide* for details).

- Single-Host Configuration (with dual path access to the module).
- Multi-Host Configuration (only support with Sun Cluster software).
- Independent Controller Configuration (no redundant path protection). Since all arrays connected to the host lose failover protection, Independent Controller Configuration should not be mixed with Single-Host Configuration and Multi-Host Configuration.

3.1.4 Tunable Parameters and Settings

- For a detailed description on tunable parameters for RAID Manager 6, see the man page `rmparams(4)`. All RAID Manager GUI processes should be terminated before changing `/etc/osa/rmparams` otherwise the GUI will core dump and exit.
- `Rdac_NoReconfig`—used to control LUN ownership between `boot -r`. Refer to the `README` file of RAID Manager 6.1.1 patch number 106552-04 or later for details. This variable is ignored by RAID Manager 6.22x.
- On hosts using volume managers, tuning `Rdac_RetryCount` to provide for quicker notification of failures is recommended since the volume manager will also do retries. A value of 1 is suggested, just as it is for clusters.
- `System_MaxLunsPerController` in `rmparams`; the higher the number, the longer it takes to reboot or invoke the GUI environment.

3.1.5 Multi-Initiator/Clustering Environment

- Sun Cluster is the only clustering/multi-initiator environment tested and verified by Sun with the A3x00. A number of parameters should be modified to run an A3x00 under the Sun Cluster 2.1/Sun Cluster 2.2 environment. See `Rdac_RetryCount`, `Rdac_NoAltOffline` and `Rdac_Fail_Flag` in the RAID Manager 6.1.1_u1 or 6.1.1_u2 patch number 106707-03 or later.

- See the Sun Cluster documentation for specific Sun Cluster requirements.

For Sun Cluster 2.1, refer to the following web site:

<http://suncluster.eng.sun.com/engineering>

For Sun Cluster 2.2, refer to the following web site:

http://suncluster.eng.sun.com/engineering/SC2.2/fcs_docs/fcs_docs.html

For Sun Cluster 3.x, refer to the following web site:

<http://suncluster.eng.sun.com/products/SC3.0/>

- Running `rm6` diagnostic commands in a multi-host configuration

When `rm6` diagnostic commands (e.g. `healthck`, `drivutil`) are run, lock files are used on the host to serialize execution and protect against these commands being run simultaneously against an A1000/A3x00/A3500FC module.

This design means that running multiple `rm6` diagnostic commands at the same time, on the same host, is safe.

Running `rm6` diagnostic commands from more than one host connected to a shared A1000/A3x00/A3500FC at the same time, however, can cause controllers to go offline, potentially affecting data availability.

This situation can occur when running explorers or third party application packages (e.g. BMC Patrol) from multiple hosts with shared A1000/A3x00/A3500FC storage at the same time.

Additionally, Sun Remote Services software runs `rm6` utilities via the scripts `healthck.sh` and `drivutil.sh` (included in the package `SRSSymod`).

In all of these cases, it's important to modify procedures or configurations to ensure `rm6` diagnostic commands are never executed from more than one host attached to a shared A1000/A3x00/A3500FC at the same time.

Running the `rm6` GUI from more than one host attached to a shared A1000/A3x00/A3500FC at the same time should also be avoided. Even when the administrator has not actively executed diagnostic commands via the GUI, an idle `rm6` GUI may nevertheless be probing the A1000/A3x00/A3500FC. The customer will still be exposed to the same problem of controllers going offline.

3.1.6 Maximum LUN Support

The default setting for maximum LUN support is 8. If more than 8 LUNs are required on each A3x00, refer to "Maximum LUN Support..." in the RAID Manager 6 Release Notes for details. Also see FIN I0589 for PCI HBAs.

Note – When performing a RAID Manager upgrade, if extended LUN support is enabled, ensure that you reenable it during the upgrade as described in the Upgrade Guide.



Caution – Do not use the `add16lun.sh` script found on the RAID Manager 6.1.1 CD-ROM on a PCI machine. Details are available in FIN I0589.

3.2 LUN Creation/RAID Level

This section contains the following topics:

- Section 3.2.1, "General Information" on page 3-5
- Section 3.2.2, "LUN Numbers" on page 3-6
- Section 3.2.3, "The Use of RAID Levels" on page 3-6
- Section 3.2.4, "Cache Mirroring" on page 3-6
- Section 3.2.5, "Reconstruction Rate" on page 3-7
- Section 3.2.6, "Creation Process (Serial/Parallel) Time" on page 3-8
- Section 3.2.7, "DacStor Size (Upgrades)" on page 3-8

3.2.1 General Information

- RAID Manager 6.22x has an "immediate LUN availability" feature: a LUN becomes Optimal in a few minutes, while the rest of the initialization is being processed in the background. See Section 3.2.6, "Creation Process (Serial/Parallel) Time" on page 3-8.
- If LUN creation is scripted, see Section 3.2.6, "Creation Process (Serial/Parallel) Time" on page 3-8.
- For multi-initiator configurations, see Section 3.3, "LUN Deletion and Modification" on page 3-9.

- DacStore size is different between LUNs created under RAID Manager 6.0/RAID Manager 6.1 vs. RAID Manager 6.1.1/RAID Manager 6.22x. See Section 3.2.7, “DacStor Size (Upgrades)” on page 3-8.
- LUN creation under RAID Manager 6.22x.
 hot_add is a new command introduced in RAID Manager 6.22x and patch 106552-04 in RAID Manager 6.1.1_u1/2. It cleans up the Solaris device tree by running devfsadm (Solaris 8 and later) or the following set of commands: drvconfig, devlinks, and disks. hot_add was available in RAID Manager 6.1.1_u1/2 in /Tools/dr_hotadd.sh on the CD.

3.2.2 LUN Numbers

- LUN numbers are allocated in sequence by RAID Manager 6 starting from zero.
- The boot device has to be LUN 0.
- You must *always* have a LUN 0. Refer to FIN I0573-2.

3.2.3 The Use of RAID Levels

- See the "RAID Level" section in Chapter 2 in the *Sun StorEdge RAID Manager 6.1.1 User's Guide* or *Sun StorEdge RAID Manager 6.22 User's Guide* for various RAID levels available on the A3x00 RAID controller. The configuration GUI also gives a description of the RAID level selected during LUN creation.
- RAID 1 and RAID 5 will take advantage of the hardware RAID controller to improve performance and data availability.
- RAID 0 LUN does not offer any redundancy. If a single drive fails, all data on the LUN is lost (see the "RAID level" description in Chapter 2 in the *Sun StorEdge RAID Manager 6.1.1 User's Guide* or *Sun StorEdge RAID Manager 6.22 User's Guide*).

3.2.4 Cache Mirroring

- See the "Changing Cache Parameters" section under the "Maintenance and Tuning" chapter in the *Sun StorEdge RAID Manager 6.1.1 User's Guide* or *Sun StorEdge RAID Manager 6.22 User's Guide*.
- Write cache is always turned off for a short duration (~ 15 min.) after a reboot or reset of the A3x00 controller module until the battery is ready.
- Write cache will be disabled when the battery strength drops below 80%.
- Write cache will be off for all LUNs if the timer of the battery is greater than two years. See Section 5.2.8, “Battery Unit” on page 5-13.



Caution – If write cache is turned on in dual/active mode you must also have it mirrored. Failure to do so may result in data corruption if a controller fails.

3.2.5 Reconstruction Rate

- Reconstruction rate depends on the "Reconstruction Rate setting" and the I/O load on the module. Refer to Chapter 7 "Maintenance and Tuning" in the *Sun StorEdge RAID Manager User's Guide*.
- Under optimal conditions (system idle with no other I/Os to the A3x00), it takes about two minutes to reconstruct a 1-GB (4+1) RAID 5 LUN under RAID Manager 6.22. If the LUN is active with host I/O request, it could take up to 10 minutes to reconstruct 1 GB of storage.
- Reconstruction can start only if certain requirements have been met. If a controller has failed, no reconstruction will start until the controller is Optimal.
- Only one reconstruction can be in progress at a time for a given controller.
- The reconstruction rate can vary by 50%, depending on the model of disk drive in the trays.
- LUNs can only be reconstructed by the controllers that own them.
- The following sequence of events indicates how the LUN status changes:
 1. Initially, in a (4+1) RAID 5 configuration, LUN X is Optimal.
 2. RAID Manager 6 detects a bad drive in LUN X.
 3. The status of LUN X then changes to degraded .
If there is a GHS, LUN X status changes to reconstructing; otherwise, LUN X remains in degrade.

Note – The GHS must be of a size greater than or equal to the failed drive.

LUN X will remain in the reconstructing state until reconstruction with GHS has been completed, even though the bad disk has been replaced.

After the reconstruction with GHS has been completed, LUN status switches to Optimal. Copyback to the replaced drive starts when LUN status changes to Optimal. New data will continue to be written to the GHS until the copyback process is complete. Then the GHS returns to the GHS pool.

- For recovery under RAID Manager 6.1.1, refer to Chapter 5 in the *Sun StorEdge RAID Manager 6.1.1 User's Guide* for details.
- For recovery under RAID Manager 6.22, refer to Chapter 6 in the *Sun StorEdge RAID Manager 6.22 User's Guide* for details.

Note – Do not revive a drive if it is failed by the controller. Refer to the *Sun StorEdge RAID Manager User's Guide*.

3.2.6 Creation Process (Serial/Parallel) Time

- LUNs can be created with either CLI or the GUI. Use the GUI to create LUNs for better coordination between the various back-end utilities.
- Be sure that an Optimal LUN 0 resides on one of the controllers and an Optimal LUN exists on the other controller before you attempt parallel LUN creation. See FIN I0573-02 or later for details on the importance of LUN 0. This condition is very important in the case of LUN creation via script with CLI.
- If you must create LUNs with CLI, allow a delay of 1 to 3 minutes between each CLI (`raidutil`) process so the device path creation on the Solaris side has a chance to complete before the next one starts.
- Parallel LUN creation means that multiple LUNs are in the formatting state as reported by the GUI.
- The creation of multiple LUNs in the same drive group is processed in serial.
- A limit of four LUNs can be created in parallel in each controller. Any more than four is queued in the controller.
- RAID Manager 6.1.1 takes 3 minutes to format 1-GB of storage under optimal conditions. For example, a *10-GB RAID 5 (4+1) LUN takes 10-GB x 3 min/GB = 30 minutes*. Refer to Chapter 3 in the *Sun StorEdge RAID Manager 6.1.1 User's Guide* for details and restrictions.
- RAID Manager 6.22x takes 2 to 5 minutes for a LUN to become Optimal while the actual format is taking place in the background. Refer to Chapter 4 in the *Sun StorEdge RAID Manager 6.22 User's Guide* for details and restrictions.

3.2.7 DacStor Size (Upgrades)

Consider the following when upgrading from RAID Manager 6.0/RAID Manager 6.1 to RAID Manager 6.1.1 and later.

DacStore size is different between LUNs created under RAID Manager 6.0/6.1 vs. RAID Manager 6.1.1/6.22x. This is an issue with customers who are running RAID Manager 6.0 or RAID Manager 6.1 and want to upgrade to RAID Manager 6.1.1 or RAID Manager 6.22x. See FIN I0557-2 for details.

RAID Manager 6.22x does not support 2-MByte DacStore. It supports 40-MByte DacStore.

3.3 LUN Deletion and Modification

- See the "Guidelines for Creating or Deleting LUNs" section in the *Sun StorEdge RAID Manager 6.22 Release Notes* for details on restrictions for LUN 0 removal. Refer to FIN I0573-2 for more information regarding the serious consequences of deleting LUN 0.
- A number of new features are available to modify a LUN/drive group while the LUN/drive group is in production. See the section "Modifying Drive Groups and Logical Units" in Chapter 4 in the *Sun StorEdge RAID Manager 6.22 User's Guide* for details.

3.4 Controller and Other Settings

This section contains the following topics:

- Section 3.4.1, "NVS RAM Settings" on page 3-9
- Section 3.4.2, "Parity Check Settings" on page 3-10

3.4.1 NVS RAM Settings

See the "NVS RAM Settings" section in the appendix of the *Sun StorEdge RAID Manager 6.1.1 Installation Guide* or the *Sun StorEdge RAID Manager 6.22 Installation Guide* for more details.

- The NVS RAM specifies the configuration of the controller. Some of the parameters that the NVS RAM controls are:
 - Transfer rate between the controller and back-end SCSI enclosure.
 - Start of Day processing .
- What version of NVS RAM to be used on a particular controller depends on the disk trays and the RAID Manager 6 releases. Refer to the following file for further details:

```
/net/artemas.ebay/global/archive/StorEdge_Products/  
sonoma/nvsram/nvsram_versions
```



Caution – Modifying the NVSRAM settings via the `nvutil (1M)` command will change the behavior of the controller. Use caution when executing this command.

3.4.2 Parity Check Settings

This section contains the following topics:

- Section 3.4.2.1, “RAID Manager 6.1.1” on page 3-10
- Section 3.4.2.2, “RAID Manager 6.22x” on page 3-10
- Section 3.4.2.3, “Parity Repair” on page 3-11
- Section 3.4.2.4, “Multi-host Environment” on page 3-11

3.4.2.1 RAID Manager 6.1.1

The RAID Manager 6.1.1 default setting is to run `parityck (1M)` once a day. This setting can be modified via "Maintenance and Tuning." See the "Changing Automatic Parity Check/Repair Settings" section in Chapter 6 in the *Sun StorEdge RAID Manager 6.1.1 User's Guide*.

If the I/O subsystem is very busy, it is likely that the `parityck` process will not complete within 24 hours. In this case, a `cron` job can be created to run `parityck` once a week or once every two weeks. The goal is to run `parityck` periodically, but there should not be multiple `parityck` processes running on the same A3x00 module at any time. `parityck` takes up I/O resources in the controller.

Also refer to bug no. 4137421 which contains a script that can speed up the time necessary to run `parityck` (checks LUNs in parallel instead of serially).

3.4.2.2 RAID Manager 6.22x

The RAID Manager 6.22x default setting for `parityck (1M)` is to run once a week. A new option to `parityck` reports parity error without repair. See the man page `parityck (1M)` for details.

If `parityck (1M)` finds a mis-matched data block and parity block, it reports the mismatch (in `/var/adm/messages` and `rmlog.log`) and regenerates new parity blocks.

You should run `parityck` with the “no repair” option, which is the default setting in RAID Manager 6.22.1. See FIN I0825. This FIN describes how to override the default from the GUI, which does not work as expected.

3.4.2.3 Parity Repair

With RAID 3 and RAID 5, data blocks are assumed to be good. Parity blocks are regenerated by `parityck (1M)` with the proper options. See the man page `parityck (1M)` for further details.

3.4.2.4 Multi-host Environment

- Only 1 host in a cluster should be capable of running `parityck`.
- Each host in a box sharing environment can run `parityck`.

System Software Installation and Configuration

This chapter provides some additional information, guidelines, and tips relating to installation and configuration of system software.

This chapter contains the following sections:

- Section 4.1, “Installation” on page 4-2
- Section 4.2, “Solaris Kernel Driver” on page 4-2
- Section 4.3, “format and lad” on page 4-4
- Section 4.4, “Ghost LUNs and Ghost Devices” on page 4-5
- Section 4.5, “Device Tree Rearranged” on page 4-9
- Section 4.6, “SNMP” on page 4-11
- Section 4.7, “Interaction With Other Volume Managers” on page 4-12

4.1 Installation

This section contains the following topics:

- Section 4.1.1, “New Installation” on page 4-2
- Section 4.1.2, “All Upgrades to RAID Manager 6.22 or 6.22.1” on page 4-2

4.1.1 New Installation

- RAID Manager 6.1.1—Refer to Chapter 1 in the *Sun StorEdge RAID Manager 6.1.1 Installation and Support Guide for Solaris*.
- RAID Manager 6.22—See the “About the Installation Procedure” section in Chapter 1 in the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide for Solaris*.

4.1.2 All Upgrades to RAID Manager 6.22 or 6.22.1

Refer to the *Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide* (part number 806-7792) at:

<http://acts.ebay.sun.com/storage/A3500/RM622>

4.2 Solaris Kernel Driver

This section contains the following topics:

- Section 4.2.1, “sd_max_throttle Settings” on page 4-3
- Section 4.2.2, “Generating Additional Debug Information” on page 4-3

Refer to the patch matrix (OS, driver, RAID Manager 6) outlined in Early Notifier 20029 for details. Also refer to Patchpro:

<http://patchpro.ebay/servlet/com.sun.patchpro.servlet.PatchProServlet>

RAID Manager 6.1.1 only supports SCSI interconnects (sd via SBUS/UDWIS/isp and PCI/glm).

RAID Manager 6.22 supports both SCSI and Fibre Channel interconnect to the A3x00/A3500FC controller module. RAID Manager 6.22 supports only SCSI when you are using the Solaris 2.5.1 11/97 operating environment. See FIN I0688. The SCSI driver stack is the same as RAID Manager 6.1.1. The driver stack for Fibre Channel is:

- SBus/soc+socal/sf/ssd
- PCI/QLC2100/ifp/ssd
- PCI/QLC220x/fcp/ssd

4.2.1 sd_max_throttle Settings

- `sd_max_throttle` for the A3x00 is set by `sd`. Manually setting `sd_max_throttle` in `/etc/system` is not necessary unless the current value is too high. A good estimate is `sd_max_throttle` x no. of LUNs per module \leq 180 for systems using an UDWIS/SBus host bus adapter (HBA). 180 is close to the upper limit of command entries in the current generation of UDWIS/SBus (SCSI) HBAs. `sd_max_throttle` has no effect on the current `ssd`.
- Refer to FIN I0579-1 for further information.

4.2.2 Generating Additional Debug Information

After setting `sd_error_level = 0` or `ssd_error_level = 0`, the following error messages may appear on the console or in `/var/adm/messages`:

```
Failed CDB:0xbe 0x1e 0x12 0x1 0xc1 0x0 0x10 0x0 0x0 0x0 0x0 0x0
/pci@1f,4000/scsi@4,1/sd@5,0 (sd50):
Sense Data:0x70 0x0 0x5 0x0 0x0 0x0 0x28 0x0 0x0 0x0 0x24
0x0 0x
0 0x0 0x0 0x0 0xe 0x12
```

The messages are warnings that the target driver `sd` or `ssd` has tried a command that didn't work. Some versions, for example, try to read the wrong mode sense page for power management data. You can safely ignore these messages.

Bug 4358075 describes error messages when the RAID Manager 6 software probes all array devices mode page 2C and receives errors from those arrays that don't support page 2C. The bug 4358075 messages are expected and are not seen unless `sd_error_level = 0` or `ssd_error_level = 0`.

The following kernel variable could be set to capture more information regarding why the controller failover occurred:

```
sd_error_level = 2 or 0 (for A3x00 SCSI)
ssd_error_level = 2 or 0 (for A3500FC)
RdacDebug = 1 (for both SCSI and FC)
```

You can set the variables in two ways:

- `adb -kw`
- You can add the variables to the end of `/etc/system`, followed by a reboot. See the man page `system (4)` for further details.

With the variables set, all failed command descriptor block (CDB) and retry commands will appear on the console and in `/var/adm/messages`. Be sure enough space is available on `/var/adm`, if file system size is limited.

4.3 format and lad

Keeping the `lad/RM6 ctd` assignment in sync with `format` is neither necessary nor practical.

`lad/RM6` reports the current device path to LUNs. The path changes when LUNs are moved between controllers.

`format` reports the device path based on entries in `/dev/[r]disk`. These entries are created at the end of the LUN creation process. These are static, thus presenting a fixed reference to the application no matter which controller owns the LUN.

If you want to keep the `ctd#` in sync, refer to the following white paper available on the Sonoma Engineering web site for a RAID Manager 6.1 solution:

```
http://webhome.sfbay/A3x00/Sonoma/4084293
```

For a RAID Manager 6.22x solution for keeping the `ctd#` in sync, see the man page `rdac_address (4)`.

Note – `format` should report the LUNs as a pseudo device (for example `/pseudo/rdnexus...`). If the path indicates physical devices (for example `/sbus@...`), then this indicates that the LUNs were not built properly.

4.3.1 Volume Labeling

At the end of the LUN creation process, `format (1M)` is called to label the LUN with a volume label. If the LUN creation process is interrupted or the LUN is created via the serial port, a valid Solaris label may not exist on the LUN. In this case, just label the LUN manually using the `format (1M)` command.

4.4 Ghost LUNs and Ghost Devices

The following sample procedure corrects a configuration with a LUN that has a drive defined at location [15,15] (not valid). The drive is an Optimal drive/LUN, and the device appears as an extra Global Hot Spare (GHS).



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

Note – This serial port procedure should be performed from the controller that owns the LUN or a reboot will be necessary. Perform a `ghsList` for information about the hot spare. You must extract the `dev` pointer information for use in subsequent steps. Make a note of the output from `ghsList`.

You must obtain the devnum of the ghost drive. Use `vdShow <LUN#>`, where `<LUN#>` is the LUN that contains the extra disk. You will see one of the devnums from this LUN in the devnums of the list of GHSs.

1. Perform a `ghsList` for information about the hot spare.

You must extract the `dev` pointer information for use in subsequent steps. Make note of the output from `ghsList`.

```
-> ghsList
```

Information about two hot spares appears:

- `dev pointer=0x2b5348c` is the address of the first Ghost Hot Spare, under GHS 0.
Make a note of it.

- dev pointer=0x2b4c3ec is the address of the second GHS, under GHS1

```
GHS ENTRY 0
dev pointer=0x2b5348c (the address of the first Ghost Hot Spare)
devnum=2
state=2
status=0
flags 4000
GHS ENTRY 1
dev pointer=0x2b4c3ec
```

```
devnum 0002
state=2
status=0
flags 4000
value = 5 = 0x5
```

2. Remove the extra LUN that is part of the Global hot spare list.

Utilize shell commands on a laptop connected directly to the RS-232 port.

```
-> m 0x dev pointer address,4
```

The memory locations are displayed 4 bytes at a time.

3. Alter the third word byte two of the dev pointer information to tell the drive it is not a hot spare.

```
-> m 0x2b5348c,4 (modifying phydev pointer in memory)
02b5348c: 02b5dd7c- (hit return to get to next word)
02b53490: 00000002- (hit return to get to next word)
02b53494: 00204000-0020000 (changing 3rd long word of phydev
structure to remove the 0x4000, but keeping 1st portion)
02b53498: 00000000-. (put decimal point "." to end writing)
value = 1 = 0x1
```

4. Write the information to disk (DacStore).

```
-> isp cfgWritePhydevDef,0x dev pointer address
value = 45155160 = 0x2b10358
```

5. Modify the global hot spare list in memory.

This entry should be the dev pointer address from ghsList. Zero this location out.

```
-> m &globalHotSpare
02f09104: 02b5348c-02b4c3ec      (Put address of second GHS here)
(packaging stack after removal of invalid entry)
02f09108: 02b4c3ec-00000000      (Zero out this location)
02f0910c: 00000000- .                (end the command with
a "." [period] and a return)
value = 1 = 0x1
```

When the globalHotSpare list is modified, the entries should be packed, removing the false GHS pointer. If there is a 0 in the middle of the list, it will cause everything after the 0 to be forgotten.

6. Write the information to disk (DacStore).

Use `isp cfgSaveGHSDrives` to save the information.

```
-> isp cfgSaveGHSDrives (Save GHS drive stack to DacStore)
value = 45155160 = 0x2b10358
```

7. Verify that only one GHS is shown.

```
-> ghsList
GHS ENTRY 0
dev pointer=0x2b4c3ec
devnum 0002
state=2
status=0
flags 4000
value = 5 = 0x5
```

4.4.1 Removing Ghost Drives

Use this procedure to remove ghost drives that are not Global hot spares. To remove a phantom drive, perform the following steps through the controller shell.



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

Note – This serial port procedure should be performed from the controller that owns the LUN or a reboot may be necessary.

1. Enter the string:

```
-> cfgPhy ch,id
```

2. Write down the `nextphy` value.

3. Enter the string:

```
-> d &phyunits,6,4
```

A column of data is returned. Each column of data following the colon (:) equals one channel.

4. Determine the channel that corresponds to the channel of the phantom drive.

The channels in the shell are 0-relative.

5. Enter the string:

```
-> d 0x(ADDRESS from step 4),0x30,4
```

6. Locate the first column that contains numbers after the colon (:).

- If the `nextphy` value copied in step 2 is shown, proceed to step 7.
- If the `nextphy` value copied in step 2 is not shown, then repeat steps 5 and 6 with the value in the first column that contains numbers after the colon (:).

7. Enter the string:

```
-> m 0x(ADDRESS from step 4),4
```

8. After the dash (-), enter the `nextphy` value from step 2 and press enter.

9. Enter a period (.) and press enter.

10. Enter the string:

```
-> cfgPhy ch,id
```

Verify “number of phydevs = 0”.

11. Enter the string:

```
-> isp cfgSaveFailedDrives
```

4.5 Device Tree Rearranged

This section contains the following topics:

- Section 4.5.1, “Dynamic Reconfiguration Related Problems” on page 4-10
- Section 4.5.1.1, “Workaround” on page 4-10

Solaris device numbers for controllers, the `c` number in `/dev/[r]dsk/cXtYdZs0`, are assigned by Solaris during reconfiguration reboots or device addition. It is not always easy to predict how and when device numbers will be re-allocated after a reconfiguration boot is done following configuration changes.

If this happens unexpectedly, mount points in `/etc/vfstab` and in Volume Manager can be lost. The file `/etc/path_to_inst` is intended to keep device numbers the same across normal reboots. Bug 4118532 describes numerous details, but note the following:

- Do not start any software installation of RAID Manager 6 or Solaris without first making sure that the device names are stable. See FIN I0727.

A reconfiguration reboot is the best way to do this. Device numbering is likely to change if hardware has been added, removed or disabled. Prior to Solaris 7, if one did a reconfiguration boot with a failed controller, the `C` number was lost because

the `disks (1M)` program would remove such failed controllers. In Solaris 7 and later, `disks` won't purge failed devices unless they are called as `disks -C` or `devfsadm -C` (Solaris 8).

- With the change to `devfsadm` in Solaris 7 and later, the file `/dev/cfg` also keeps bus numbers and may need to be removed before a reconfiguration reboot in order to clear up persistent misnumbering.

4.5.1 Dynamic Reconfiguration Related Problems

The first time you add an array, the `/kernel/drv/rdriver.conf` file serves much the same purpose as the `/kernel/drv/sd.conf` file, except that the `rdriver` takes the place of `sd`. The `rdriver` reads this file and determines the number of LUNs that it will allow to be configured in the `/devices/pseudo/rdnexus@?` device tree. This is the same device tree where links are created from `/dev/[r]dsk`.

The `/kernel/drv/rdriver.conf` file contains two types of entries.

- Specific LUN definitions that match each configured LUN. They appear at the top of the file.
- The second type resemble those found in the `/kernel/drv/sd.conf` file and occur at the bottom of the file. These entries show only LUNs 0-7 in the default `rdriver.conf` file.

The problem, then, is that when we DR Attach the A3500's, the specific or actual LUN definitions are not placed into the `/kernel/drv/rdriver.conf` file until the `dr_hotadd.sh` (RAID Manager 6.1.1) or `hot_add` (RAID Manager 6.22x) script is run. This is too late: the `rdriver` has already been loaded during the last reboot, and `rdriver` has already read the `rdriver.conf` file, which did not contain the new, higher than LUN 8 A3500 entries. So the new, higher than 8 LUNs will not be recognized without a reboot.

4.5.1.1 Workaround

Modify the entries at the bottom of the `rdriver.conf` file to allow targets 4 and 5 (or whichever targets your system uses) to accept more than 8 LUNs: 16 LUNs, for example. You must reboot the system at least once to make this effective. However, after the reboot, the `rdriver` is loaded and ready to accept LUNs greater than 8 dynamically. You can then attach one or more A3500s, and when the `dr_hotadd.sh` (RAID Manager 6.1.1) or `hot_add.sh` (RAID Manager 6.22) script is run, the `rdriver` is prepared to allow the creation of LUNs greater than 8, with the appropriate device entries.

- Once you determine the desired maximum number of LUNs you want, complete the configuration by running the `addXXlun.sh` from the Tools directory of the CD-ROM. An online version of the Tools directory is available from:

```
/net/artemas.ebay/global/archive/StorEdge_Products/  
sonoma/rm_6.1.1_u2/FCS/Tools
```

or

```
/net/artemas.ebay/global/archive/StorEdge_Products/  
sonoma/rm_6.22/Tools
```

Note – Adding support of more LUNs than you need extends the time required for reboot and the response time of the RAID Manager 6 GUI because it has to scan all the potential LUNs. See FIN I0551-1 or later.

- Common problem—RAID Manager 6 is unable to communicate with the module but `lad` shows more than 8 LUNs.

Solution—Re-run `addXXlun.sh` followed by `boot -r` or `hot_add`.

4.6 SNMP

See the "Setting Up SNMP Notification" section in Chapter 4 in the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide* for details.

Note – In order for SNMP to be properly configured, DNS must be enabled unless you apply the workaround described in bug 4348634.

The SNMP trap data that is supported by the A3x00/A3500FC controllers is covered in the MIB definition included with the RAID Manager host software.

See the following file:

```
/net/artemas.ebay/global/archive/StorEdge_Products/  
sonoma/rm_6.22/rm6_22_FCS/Product/SUNWosau/reloc/lib/osa/  
rm6traps.mib
```

Also refer to the RAID Manager software documentation online:

```
http://infoserver.central
```

4.7 Interaction With Other Volume Managers

This section contains the following topics:

- Section 4.7.1, “VERITAS” on page 4-12
- Section 4.7.2, “Solstice DiskSuite (SDS)” on page 4-13
- Section 4.7.3, “Sun Cluster” on page 4-13
- Section 4.7.4, “High Availability (HA)” on page 4-13
- Section 4.7.5, “Quorum Device” on page 4-14

4.7.1 VERITAS

This section contains the following topics:

- Section 4.7.1.1, “VERITAS Enabling and Disabling DMP” on page 4-12
- Section 4.7.1.2, “HA Configuration Using VERITAS” on page 4-13
- Section 4.7.1.3, “Adding or Moving Arrays Under VERITAS” on page 4-13

The A3x00 has been tested and qualified with VERITAS Volume Manager (VxVM) and SDS. Both VxVM and SDS are layered software running on top of `rdriver`.

The only A3x00 cluster environment that has been tested and qualified by Sun is Sun Cluster 2.1 and Sun Cluster 2.2. Other cluster software (VERITAS Cluster Server and FirstWatch, for example) has not been tested nor qualified by Sun.

4.7.1.1 VERITAS Enabling and Disabling DMP

- The following document, available on the Sonoma Engineering web site, provides instructions for installing, running, and administrating VxVM with A3x00:

http://webhome.sfbay/A3x00/Sonoma/VM_A3x00_A1000.pdf

- VERITAS VxVM supports DMP (Dynamic Multipathing). This feature conflicts with the dual path control of RAID Manager 6. VxVM release 3.0.2 addressed this issue so that RAID Manager 6.22x can coexist with DMP on the same host. DMP should be disabled under previous versions of VxVM. Refer to the *Sun StorEdge RAID Manager 6.22 Release Notes* for instructions on disabling DMP.

4.7.1.2 HA Configuration Using VERITAS

If you have a problem running an A3x00 under a third party cluster environment, you can check with CPRE to see whether they have VIP arrangement with the third party vendor to help you move forward.

Because Sun has no access to the source code of third party cluster software, debugging is problematical. In such a case, a SCSI or Fibre Channel analyzer trace between the host and the A3x00 module would help to isolate whether the issue is in the upper software layer or in the A3x00 layer.

4.7.1.3 Adding or Moving Arrays Under VERITAS

Under VERITAS, volumes will show up in the Vxdisk list as *online altused* after being moved to a new host. See SRDB no. 20907 for further details.

4.7.2 Solstice DiskSuite (SDS)

The A3x00 is supported by Solstice DiskSuite 4.1. See the following white paper available on the Sonoma Engineering web site:

<http://webhome.sfbay/A3x00/Sonoma/SDS.ps>

4.7.3 Sun Cluster

Refer to Section 2.6.1, "Cluster Information" on page 2-10.

4.7.4 High Availability (HA)

High Availability and Parallel Database were two distinct products offered by SunSoft and SMCC. The functionality of these two products is available in Sun Cluster 2.1. See Section 1.4.1 "High Availability and Parallel Database Configurations" in Chapter 1. The document is available on the following web site:

http://suncluster.eng/engineering/SC2.2/fcs_docs/fcs_docs.html

4.7.5 Quorum Device

Quorum is a concept that is used in distributed systems, a cluster environment particularly. The requirements and restrictions of a quorum device are specific to the particular cluster environment. Refer to the following web sites for online documentation:

<http://suncluster.eng.sun.com/engineering/SC2.1>

http://suncluster.eng/engineering/SC2.2/fcs_docs/fcs_docs.html

Using the Sun StorEdge A1000 or A3x00/A3500FC array as a quorum device is not supported. See FIN I0520-02.

Maintenance and Service

This chapter provides maintenance and service information for verifying FRU functionality, guidelines for replacing FRUs, and tips on upgrading to the latest software and firmware levels.

This chapter contains the following sections:

- Section 5.1, “Verifying FRU Functionality” on page 5-2
- Section 5.2, “FRU Replacement” on page 5-10
- Section 5.3, “Software and Firmware Guidelines” on page 5-16

5.1 Verifying FRU Functionality

This section contains the following topics:

- Section 5.1.1, “Disk Drives” on page 5-3
- Section 5.1.2, “Disk Tray” on page 5-4
- Section 5.1.3, “Power Sequencer” on page 5-5
- Section 5.1.4, “SCSI Cables” on page 5-6
- Section 5.1.5, “SCSI ID Jumper Settings” on page 5-7
- Section 5.1.6, “SCSI Termination Power Jumpers” on page 5-7
- Section 5.1.7, “LED Indicators” on page 5-7
- Section 5.1.8, “Backplane Assembly” on page 5-7
- Section 5.1.9, “D1000 FRUs” on page 5-7
- Section 5.1.10, “Verifying the HBA” on page 5-8
- Section 5.1.11, “Verifying the Controller Boards and Paths to the A3x00/A3500FC” on page 5-8
- Section 5.1.12, “Controller Board LEDs” on page 5-9
- Section 5.1.13, “Ethernet Port” on page 5-10

The troubleshooting and replacement procedures for the following controller module FRUs are documented in detail in the *Sun StorEdge A3500/A3500FC Controller Module Guide*:

- Battery Canister F370-2434
- Controller Fan F370-2433
- DC Power and Battery Harnesses F565-1397
- Power Supply F370-2436
- Power Supply Fan F370-2432
- Power Supply Housing F370-2869
- Mounting Rail 370-3655

Note – If a fan failure message appears in the `rmlog.log`, replace the fan FRU that was reported to have failed (controller or power supply fan) even if the fan appears to be spinning. The fan circuitry has an RPM sensor which triggers the fan failure message. The fan may continue to spin but in a degraded mode (at half speed).

Note – Remember to reset the battery date on both controllers after a battery replacement. Refer to Chapter 6 in the *Sun StorEdge RAID Manager 6.22 User's Guide* and read the section “Recovering from Battery Failures” for details on resetting the battery date.

Note – The power supplies have a thermal protection shutdown feature. To recover from a power supply shutdown, see Section 7.1 “Recovering From a Power Supply Shutdown” in the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

5.1.1 Disk Drives

Refer to Chapter 3 in the *Sun StorEdge RAID Manager 6.22 User's Guide* for procedures to verify the status of each disk drive.

When a disk drive fails, the disk drive amber LED should be on. See Section 4.3.7 “Disk Drive Problem” in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for further details.

Note – RAID Manager Recovery Guru will lead you through a disk drive replacement step by step. Do not deviate from this procedure or you may end up with ghost drives, ghost LUNs, or *drive has not been detected* problems. For the same reason, do not swap disk drives while the controller module is powered off.

Make sure each disk drive has the supported firmware level. If there is a failure affecting the entire disk tray, RAID Manager 6 Recovery Guru will report the drive side channel failure. You need to first fix the drive side channel failure before any drive can be reconstructed (see Section 5.2.10, “Disk Drives” on page 5-14).

Note – Do not swap drives between an A1000 and an A3x00 or you may end up with DacStore and NVSRAM corruption.

Replacing a failed drive in a RAID 0 LUN requires special attention. Stop any volume manager or upper level software from accessing the LUN to prevent possible data corruption. As soon as the drive is replaced, the LUN is described as “optimal.” Because the data was lost since it is a RAID 0 LUN, you must reformat the LUN. See page 10 of the *Sun StorEdge RAID Manager 6.22.1 Release Notes* for more information.

Never “revive” drives unless a LUN is optimal, as described in “Recovery Guru Revive Option Is Removed” in the *Sun StorEdge RAID Manager 6.22.1 Release Notes*.

5.1.2 Disk Tray

This section contains the following topics:

- Section 5.1.2.1, “RSM Tray” on page 5-5
- Section 5.1.2.2, “D1000 Tray” on page 5-5

Several conditions can cause a disk tray to become inaccessible: a loose or defective SCSI cable, a loose or defective SCSI terminator, a defective SCSI chip on the controller board, or a defective component in the disk tray.

The problem can be sometimes difficult to isolate. Check the `rmlog.log` and system logs for an error sense code or a FRU code. Check each cable connector and look for bent pins. Ensure that each cable is properly connected.

A disk tray failure can cause all of the drives to report a failed status. RAID Manager 6 Recovery Guru can be used to determine the drive side channel failure. Individual drives are not recoverable until the drive side channel failure status has been resolved.

To perform further troubleshooting, you will need to have spare components on hand: a controller FRU, SCSI cables, SCSI terminators, and the disk tray interface board.

After the hardware component has been replaced, run a health check to verify the status of the drive channel, drive tray, and each disk drive. You might need to power cycle the controller module (not the entire rack), even though RAID Manager 6 Recover Guru instructions indicate that you don't have to power cycle the controller module if the controller firmware level is 2.5.2 or higher.

You need to power cycle the controller module if any of the following conditions occur:

- The LUN reconstruction does not start.

Note – LUN reconstructions occur one at a time serially.

- Health check continues to report drive side channel failure.
- Any of the drives in the disk tray do not come out of the failed or unresponsive state to reconstruct.

Power cycling of the controller module enables the controller to re-scan the disk trays. If you continue to encounter error messages while attempting to bring a drive back online (reconstruct), physically remove the drive, wait 30 seconds, and then reinstall the drive. Then use RAID Manager 6 Recovery Guru to reconstruct the drive. See FIN I0670 about replacing the ESM card.

5.1.2.1 RSM Tray

A common point of failure on the RSM tray is with the WD2S card, part number 370-2196 (older version) and part number 370-3375 (newer version). This card is located at the point where the SCSI cable attaches to the RSM disk tray. The function of this card is to convert Wide Differential SCSI to Single Ended SCSI.

The other common point of failure on the RSM tray is with the SEN card, part number 370-2195. The SEN card should have microcode rev 1.1. When you replace the Wide Differential 2 SCSI card, change the SCSI ID range from 0-6 (default) to 8-14, which is needed to see the SCSI IDs properly in a Sun StorEdge A3x00 array. Failure to do so results in a ghost drive [15,15].

5.1.2.2 D1000 Tray

The common point of failure on the D1000 tray is with the differential SCSI controller board, part number 375-0008.

The disk tray midplane is also a single point of failure but only in rare cases. In very rare cases a failed drive can tie up the entire SCSI bus causing the remaining drives to be rendered inaccessible. This problem can be easily mistaken as a bad disk tray and somewhat difficult to isolate. Compare each drive's LED pattern. In some cases a failed drive will indicate a different LED pattern from the remaining drives.

5.1.3 Power Sequencer

The power sequencer has two groups of four sequenced power outlets and one unsequenced power outlet. They are designated sequenced group 1 and sequenced group 2. There is a four second power on delay between sequenced group 1 and sequenced group 2. The unsequenced outlet remains on as long as the sequencer circuit breaker is turned on and the sequencer is connected to an AC power source. When one half of the rack does not receive power, it is usually the result of power sequencer failure.

The first thing to do is to verify that the power sequencer is plugged in to a 220 VAC power source.

Note – At the bottom of the expansion cabinet are two power sequencers. The front power sequencer is hidden behind the front key switch panel. Remove the front key switch panel to gain access to the power sequencer's power cable.

There is a Local/Remote switch located on the front panel of each power sequencer. When the Local/Remote switch is set to Local the sequenced outputs are controlled by a circuit breaker located on the front panel of each power sequencer. When the Local/Remote switch is set to Remote, the sequenced outputs are controlled by the key switch located at the bottom front of the Expansion rack. When the Local/Remote switch is set to OFF, power is removed from the sequenced outputs. The unsequenced output is not affected by the position of the Local/Remote switch. As long as the AC power cord is connected and AC power is available, the unsequenced output is available.

One possible failure is that only one of the sequenced groups turns on. In this case you will see only one power supply functioning in some of the disk trays installed in the Expansion rack. Refer to Chapter 3 in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide* and verify that each disk tray is properly configured.

In a 3x15 configuration, the power sequencers in each rack need to be daisy chained front to front and back to back (See Section 3.5.2 “Connections Between Power Sequencers” in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*). The Local/Remote switch on each power sequencer should be set to *Remote*. The 2x7 rack has a front key switch that controls both racks. The 1x8 rack does not have a front key switch.

5.1.4 SCSI Cables

The most common problem involving SCSI cables is with bent pins on the connectors. This usually occurs during a system installation. A typical indication of a defective SCSI cable is an error message indicating SCSI parity errors and the SCSI transfer rate has been reduced or has switched from wide to narrow SCSI. A SCSI cable that has failed on the host side usually results in a data path failure indication. If the problem is with a failed SCSI cable on the drive side, the result is usually in the form of a drive side channel failure indication. Currently there are no procedures for testing a SCSI cable for failure other than to replace it with a new or known good cable. Also, ensure that the SCSI bus length is within the recommended maximum of 25 M (see Section 2.1.5, “SCSI and Fiber-Optic Cables” on page 2-3).

5.1.5 SCSI ID Jumper Settings

The controller module SCSI ID can be changed, if necessary, by the use of jumpers. The SCSI ID jumper block is located at the rear of the SCSI controller module. See Section 2.3 “Verifying Controller Module ID Settings” in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide* for detailed instructions. The factory default settings are shown in TABLE 5-1.

TABLE 5-1 Controller Module SCSI ID Settings

Controller	SCSI ID
A	5
B	4

5.1.6 SCSI Termination Power Jumpers

There are two SCSI termination power jumpers located at the rear of the SCSI controller module. One is located below and to the left of the J11 connector (DIFF SCSI ARRAY 1). The other is located below and to the right of the J4 connector (DIFF SCSI HOST B). These jumpers enable the controller boards to supply SCSI termination power to the host SCSI bus. Do not remove these jumpers.

5.1.7 LED Indicators

See Section 4.1 “Checking the Controller Module LEDs” in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for detailed information on checking the controller module LEDs.

5.1.8 Backplane Assembly

Backplane failures are rare. There are no active devices on the backplane.

5.1.9 D1000 FRUs

Refer to the following web site for a listing of the D1000 FRUs:

<http://infoserver.central/data/syshbk>

5.1.10 Verifying the HBA

- Refer to Early Notifier 20029 for the latest information regarding HBA support.
- The UDWis/SBus host bus adapter (HBA) should be at firmware level 1.28 or higher (Refer to FCO A0163-1 and FIN I0547 for further details).
- The older SOC+ card, part number 501-3060, is not supported with the A3500FC. You need to check the card part number label located on the SBus connector to determine the part number. Do not rely on the output of `prtdiag`, it may provide the wrong part number.
- The newer SOC+ card 501-5202 and 501-5266 are supported with A3500FC. Be sure to have firmware level 1.13 (patch no. 109400-03 or higher) installed.
- To use the onboard SOC+ you need to have firmware level 1.13 installed (patch no. 103346-25 or higher). The two onboard SOC+ on one I/O board can be used at the same time.

5.1.11 Verifying the Controller Boards and Paths to the A3x00/A3500FC

See Section 4.1.2 “Checking the Controller LEDs” in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for information on interpreting the controller status LED pattern.

A controller held in reset (offline) does not necessarily mean that the controller is defective. Rather, it indicates that the I/O between the controller and the host has been interrupted. There are several conditions that can cause this: a defective I/O board, a defective HBA, a defective host SCSI cable or bent pins on the cable, a defective SCSI terminator, a defective controller (either one), or the user may have taken the controller offline manually.

RAID Manager 6 Recovery Guru provides step by step instructions to assist you in troubleshooting a data path failure or an offline controller. If you have already replaced the controller board and the replacement controller still cannot be brought online, this is a good indication that the data path failure is with another defective component.

The following two commands are very helpful in troubleshooting data path failures: `rdacutil -U` and `rdacutil -u`.

The `rdacutil -U` command unfails the alternate controller without sending I/O through the data path to check the controller. The controller goes through SOD self diagnostics which is not an extensive diagnostics test. If the controller passes SOD then it should come online.

The `rdacutil -u` command unfails the alternate controller then attempts to communicate with the alternate controller through the I/O path. This is how the RAID Manager 6 GUI unfails a controller.

Sometimes by issuing these two commands you can determine if the failure is internal to the controller board or external. If the controller is not at fault, to perform further troubleshooting, you will need to have spare parts readily available to isolate the defective component by substituting components one at a time. A defective HBA is one possible cause. Verify that the firmware used in the HBA (for example firmware 1.28 or higher for an UDWIS/SBus HBA) is current. A defective host SCSI cable or terminator is another possible cause for data path failure. Sometimes you may see *reducing sync transfer rate* or *disabled wide SCSI mode* error messages. The most common cause for these error messages are bent pins on a SCSI cable connector. Also be aware that a SCSI terminator can be defective even if the green LED on the SCSI terminator is on.

Note – Do not perform a `boot -r` while a controller is held in reset because it will rearrange the device path and may cause the controller to not appear in RAID Manager 6.

5.1.12 Controller Board LEDs

The controller module's LEDs indicate the status of both the controller module and its individual components. The green LEDs indicate normal operating status; amber LEDs indicate a hardware fault. It is important that you check all the LEDs on the front and back of the controller module when you turn on power. Besides fault isolation, the LEDs can be used to determine if there is any I/O activity between the host and the controller modules.

Here are a few things to consider when checking the LEDs for status:

- If a Fast Write Cache operation or other I/O activity is in progress to the controller module (or attached drive units), you may see several green LEDs blinking including: the Fast Write Cache LED (on the front panel), controller FRU status LEDs, or applicable drive activity LEDs.
- The green *heartbeat* LED on the controller FRUs blinks continuously.
- See Section 4.1.2 “Checking the Controller LEDs” in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for the LED pattern information.
- An active controller will not have the same status LEDs lit as a passive controller.
- If you just turned on power, the controller module's green and amber LEDs will turn on and off intermittently. Wait until the controller module finishes powering up before you begin checking for faults.

5.1.13 Ethernet Port

The Ethernet port located on the back of the controller module is not supported.

5.2 FRU Replacement

This section contains the following topics:

- Section 5.2.1, “HBA” on page 5-10
- Section 5.2.2, “Interconnect Cables” on page 5-11
- Section 5.2.3, “Power Cords” on page 5-11
- Section 5.2.4, “Power Sequencer” on page 5-11
- Section 5.2.5, “Hub” on page 5-12
- Section 5.2.6, “Controller Card Guidelines” on page 5-12
- Section 5.2.7, “Amount of Cache” on page 5-13
- Section 5.2.8, “Battery Unit” on page 5-13
- Section 5.2.9, “Cooling” on page 5-14
- Section 5.2.10, “Disk Drives” on page 5-14
- Section 5.2.11, “Disk Tray” on page 5-14
- Section 5.2.12, “Midplanes” on page 5-15
- Section 5.2.13, “Reset Configuration and `sysWipe`” on page 5-16

5.2.1 HBA

If the host server does not support hot swapping of the I/O boards, you will need to shut down the host to replace an HBA. You should read the manual that comes with the HBA and become familiar with the HBA installation procedures.

Note – There are several sections in the *Sun Enterprise Cluster System Hardware Service Manual* that provide detailed procedures on how to disconnect a host and to remove and replace an HBA or a RAID controller.

5.2.2 Interconnect Cables

- Host SCSI cables

Stop all I/O activities to the corresponding data path before replacing a host SCSI cable.
- SCSI cables

See Section 6.1 in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for further details.
- Fiber-optic cables

See Section 6.2 in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for further details.
- Controller module terminators

Stop all I/O activities to the corresponding controller module before replacing a controller module terminator.
- Drive side SCSI cables

You should stop all I/O activities to the controller module and power down the controller module before replacing the drive side SCSI cables. If the failure caused a disk tray to lose communication to the controllers, you will need to power cycle the controller module to re-establish communication with the disk tray.

5.2.3 Power Cords

- Power cords from the power sequencer to the device trays

See Section 7.2 in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for further details. You will need to power down the entire rack to perform this operation. You will need to remove the bottom disk tray to gain access to the power sequencer outlets. Further, you will need to remove the rack side panel to perform this service.
- Power cords from the power sequencer to the AC source.

See Section 5.3 in the *Sun StorEdge Expansion Cabinet Installation and Service Manual* for further details.

5.2.4 Power Sequencer

See Section 5.4 in the *Sun StorEdge Expansion Cabinet Installation and Service Manual* for further details. You will need to power off the entire rack to replace a power sequencer. Be sure each AC cords are plugged back to its original outlet position.

5.2.5 Hub

You need to stop all I/O activities to the hub before replacing it. Refer to the *FC-100 Hub Installation and Service Manual* for further details.

5.2.6 Controller Card Guidelines

- With RAID Manager 6.22.1 or patches 109232 and 109233, there are new NVSRAMs. With a Sun StorEdge A1000, download the NVSRAM after the controller card is replaced. See FIN I0709.

Note – Because controller FRUs come with an NVSRAM, remember to always hot-swap A3x00/A3500FC controller FRUs to preserve the NVSRAM that exists on the failed controller. See FIN I0709.

- To replace a failed controller board, see Section 6.3.1 in the *Sun StorEdge A3500/A3500FC Controller Module Guide* and refer to the *Sun StorEdge A3x00 Controller FRU Replacement Guide*. Also refer to FIN I0553-1.
- Refer to Section 2.6.1, “Cluster Information” on page 2-10 for guidelines when replacing a controller within a cluster environment.

Note – SCSI controller FRUs are factory loaded with a universal firmware that can be upgraded or downgraded. Make sure to follow the procedures documented in the *Sun StorEdge A3x00 Controller FRU Replacement Guide* to download firmware to the newly replaced controller to match with the supported RAID Manager on the host. Firmware downgrades on controllers that are not universal FRUs can only be done through the serial port.



Caution – Do not use RAID Manager GUI to downgrade the firmware. Use the command line to downgrade the firmware.



Caution – The possibility of controller “deadlock” exists with certain A3x00 and RAID Manager 6.1.x configurations. Refer to FIN I0643-01 prior to performing any firmware upgrade.

For a universal FRU firmware downgrade, you need to load appware first then bootware. For a firmware upgrade, you need to load bootware first then appware.

There are four different controller FRUs:

- SCSI Controller Canister w/Memory (D1000), part number 540-3083

- SCSI Controller Canister w/Memory (RSM), part number 540-3600
- Fiber Controller Canister w/Memory (D1000), part number 540-4026
- Fiber Controller Canister w/Memory (RSM), part number 540-4027

When returning a controller canister for repair, ensure that the memory SIMMs are returned with the controller canister. If a SCSI controller canister being returned has 128-MB of cache memory, order two memory FRUs, part number F370-2439 in addition to the replacement controller canister FRU.

Note – You might sometimes come across a controller with part number 375-0089-01. This controller has the TX CDC chip. It was shipped in the SCSI version of the A3500 with D1000 disk trays from July 1999 through October 1999. This controller will only work with the SCSI versions of the A3500 (with D1000 disk trays) systems that have RAID Manager 6.1.1 Update 2 or higher. The replacement FRU for this controller is part number 540-3083. ***The TX controller should not be returned to the FRU supply inventory.***

5.2.7 Amount of Cache

- The SCSI controller FRU is configured at the factory (default) with 64 MBs of cache memory per board.
- The SCSI controller FRU cache memory can be upgraded in the field to 128 MBs.
- The FC controller FRU is configured at the factory (default) with 128 MBs of cache memory per board.
- Write cache in the A3x00 is normally mirrored between two controllers thus reducing the effective cache size available to a controller by half.
- The amount of cache memory between two controllers has to match before the *Write Cache Mirroring* parameters described in Chapter 7 in the *Sun StorEdge RAID Manager 6.22 User's Guide* can be changed.

5.2.8 Battery Unit

See Section 7.4 “Replacing the Battery Unit” in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for detailed instructions on replacing the battery unit.

- Make sure you reset the battery age on both controllers.
- Check the battery date code label to make sure the new battery FRU has not exceeded its shelf life (12 months from date of Manufacture)

- The battery has a service life of two years. After two years, it needs to be replaced. A fresh battery will guarantee that the data saved in the controller's cache memory will be kept live for up to the design specification of 72 hours.
- See the sections "To Replace Old Batteries" and "To Replace New Batteries" in the *Sun StorEdge RAID Manager 6.22.1 Release Notes* for more information on replacing old and new batteries.

5.2.9 Cooling

See Section 7.1 "Recovering From a Power Supply Thermal Shutdown" in the *Sun StorEdge A3500/A3500FC Controller Module Guide* for further steps to take in the event of a cooling related problem.

5.2.10 Disk Drives

Do not install disk drives from other arrays into an A3x00/A3500FC (for example, A1000, A3000, etc). This can cause DacStore corruption in the array and will require another download of the NVSRAM file to repair. For further information regarding NVSRAM, refer to the `nvsram_versions` file in the following internal directory:

```
/net/artemas.ebay/global/archive/StorEdge_Products  
/sonoma/nvsram/nvsram_versions
```

When you replace a disk drive, if it is not already in the failed state, make sure you use RAID Manager 6 to fail it before removing the disk drive. Then replace the disk drive and use RAID Manager 6 GUI to bring the new disk drive into the configuration.

Swapping drives without first using RAID Manager 6 to fail it or removing and replacing a drive when the controller is powered off can result in phantom disk drives or phantom LUNs. Further, it will require serial port access to remove the phantom drive. This applies to the Global Hot Spare (GHS) drive as well.

5.2.11 Disk Tray

Refer to Chapter 3 in the *Sun StorEdge A1000/D1000 Installation, Operations and Service Manual* for detailed instructions on removing and replacing components.

Power to the disk tray needs to be turned off while replacing the following disk tray components:

- The D1000 interface board (see FIN I0670-1 for proper procedure)

- In the RSM tray: the SEN card and the WD2S card
Install a jumper at location ID3 to change SCSI address range from 0-7 to 8-15.
- The entire disk tray

RAID Manager 6 Recovery Guru reports drive side channel failure when the failure affects the entire disk tray. After the hardware component has been replaced, run a health check to verify the status of the drive channel, drive tray, and each disk drive. You might need to power cycle the controller module (not the entire rack), even though RAID Manager 6 Recover Guru instructions indicate that you don't have to power cycle the controller module if the controller firmware level is 2.5.2 or higher.

You need to power cycle the controller module if any of the following conditions occur:

- The LUN reconstruction does not start.

Note – LUN reconstructions occur one at a time serially per controller.

- Health check continues to report drive side channel failure.
- Any of the drives in the disk tray do not come out of the failed or unresponsive state to reconstruct.

Power cycling of the controller module enables the controller to re-scan the disk trays. If you continue to encounter error messages while attempting to bring a drive back online (reconstruct), physically remove the drive, wait 30 seconds, and then reinstall the drive. Then use RAID Manager 6 Recovery Guru to reconstruct the drive.

Note – Remember to replace any dummy drives that were removed during disk tray service. The dummy drives are important to maintain proper air flow and cooling in the disk tray.

5.2.12 Midplanes

See Section 6.3.3 in the *Sun StorEdge A3500/A3500FC Controller Module Guide* to remove and replace the controller card cage with the midplane.

5.2.13 Reset Configuration and sysWipe

Reset Configuration or `sysWipe` will delete all LUNs and bring the RAID system to a default state: active controller A, passive controller B and one default 10-MB LUN 0.

Reset Configuration is a RAID Manager 6 procedure and `sysWipe` is a serial port command. `sysWipe` wipes clean all prior DacStore data. You need to issue a `sysReboot` after a `sysWipe` command is executed.



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

Note – When you issue a `sysWipe` command you may see a message indicating that `sysWipe` is being done in a background process. Wait for a follow-on message indicating that `sysWipe` is completed before issuing a `sysReboot` command. `sysWipe` should be run from each controller.

There are times when a known good controller A is held in reset by controller B so you will not be able to access the shell tool to issue a `sysWipe` command for controller A. In this case, perform the following:

1. **Physically remove controller A and controller B.**
2. **Then insert controller A but leave controller B out temporarily.**
3. **Issue a `sysWipe` command followed by a `sysReboot` command for controller A.**
4. **Then physically remove controller A and insert controller B.**
5. **Issue a `sysWipe` command followed by a `sysReboot` command for controller B.**
6. **Then insert controller A back into the system and the system should be in a default state.**

5.3 Software and Firmware Guidelines

This section contains the following topics:

- Section 5.3.1, “Firmware, Software, and Patch Information” on page 5-17
- Section 5.3.2, “RAID Manager 6 Upgrade” on page 5-18

- Section 5.3.3, “Firmware Upgrade” on page 5-18

Follow these general guidelines to simplify software and firmware upgrades.

- Do not run an A3x00 with mixed firmware levels. FRU replacements for controller boards come at firmware 02.05.06.32. This level can be upgraded (or downgraded) to an appropriate level after installation. *Always check the firmware level.*
- When upgrading from RAID Manager 6.0 to RAID Manager 6.1.1 (or later), upgrading the firmware requires an intermediate step. *The firmware must first be upgraded to 02.04.04, and then from there may be upgraded to 02.05.02. This procedure is clearly stated in the release notes.*
- When adding newer A3x00 hardware with older A3x00 hardware, all hardware can be run with the latest software. It is best, however, to migrate older firmware to match the newer firmware on the system.
- When running RSM2000/A3000 units on same system with A1000 units, firmware in RSM2000/A3000 units must be at 02.05.02.04 or newer to maintain compatibility. A1000 units work with RAID Manager 6.1.1 and above.
- The current *Sun StorEdge RAID Manager 6.22.1 Release Notes*, are available on the following web site:

`http://infoserver.central`

- Refer to the *Sun StorEdge A1000 and A3x00 Installation Supplement* on the following web site:

`http://docs.sun.com:80/ab2/coll.266.1/ISUPPA1000A3X00/@Ab2PageView/85?`

This document contains a workaround for a disk drive firmware download bug in the A3000 only. This workaround is a script that stops SEN card polling in the A3000 controller, downloads new firmware, and then restarts SEN card polling.

5.3.1 Firmware, Software, and Patch Information

For detailed information regarding firmware, software, and patch information, refer to Early Notifier 20029 located on the following web site:

`http://sunsolve.Ebay.Sun.COM/cgi/retrieve.pl?doc=enotify%2F20029&zone_32=20029`

Also refer to the PatchPro web site:

`http://patchpro.ebay/servlet/com.sun.patchpro.servlet.PatchProServlet`

5.3.2 RAID Manager 6 Upgrade

You should use the *Sun StorEdge RAID Manager 6.22 and 6.22.1 Upgrade Guide* (part number 806-7792) to perform any upgrade to RAID Manager 6.22 or 6.22.1.

Refer to the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide for Solaris* for further details.

RAID Manager 6.1.1 does not support Solaris 8. Solaris 8 support requires RAID Manager 6.22 and patch no. 108553. For Solaris 9, only RM 6.22.1 is supported.

RAID Manager 6.0 and RAID Manager 6.1 has a 2-MB DacStore. RAID Manager 6.1.1 and higher has a 40-MB DacStore. Refer to FIN I0557.

5.3.3 Firmware Upgrade



Caution – The possibility of controller “deadlock” exists with certain A3x00 and RAID Manager 6.1.x configurations. Refer to FIN I0643-01 prior to performing any firmware upgrade.

The firmware upgrade steps are as follows (also refer to Appendix A in the *Sun StorEdge RAID Manager 6.22 Installation and Support Guide for Solaris*).

RM 6.0		RM 6.1		Universal		RM 6.22
2.4.1d	to	2.4.4.1	to	2.5.6.32	to	3.1.2.35

In each of the steps, upgrade bootware first then Appdware.

Chapter 7 “Maintenance and Tuning” in the *Sun StorEdge RAID Manager 6.22 User’s Guide* provides specific information for downloading firmware in either Online or Offline mode.

Note – The only valid time to downgrade firmware is from a universal FRU with firmware level 2.5.6.32. Any other situation requires that it be done through the serial port. Refer to FIN I0553-1 for further details.



Caution – Only trained and experienced Sun personnel should access the serial port. You should have a copy of the latest *Debug Guide*. There are certain commands that can destroy customer data or configuration information. No warning messages appear if a potentially damaging command has been executed.

Note – If you are running Solaris 7 dated 11/99 and plan to upgrade the firmware on an A3x00 controller, you need to ensure that patch no. 106541-10 (KJP 10) for Solaris has been installed. Refer to Sun Early Notifier EN20029 and bug 4334814 for further details.

Troubleshooting Common Problems

This chapter discusses some common problems encountered in the field and provides additional information and tips for troubleshooting these problems.

This chapter contains the following sections:

- Section 6.1, “Controller Held in Reset, Causes, and How to Recover” on page 6-2
- Section 6.2, “LUNs Not Seen” on page 6-6
- Section 6.3, “Rebuilding a Missing LUN Without Reinitialization” on page 6-7
- Section 6.4, “Dynamic Reconfiguration” on page 6-11
- Section 6.5, “Controller Failover and LUN Balancing Takes Too Long” on page 6-12
- Section 6.6, “GUI Hang” on page 6-13
- Section 6.7, “Drive Spin Up Failure, Drive Related Problems” on page 6-13
- Section 6.8, “Phantom Controllers Under RAID Manager 6.22” on page 6-14
- Section 6.9, “Boot Delay (Why Booting Takes So Long)” on page 6-15
- Section 6.10, “Data Corruption and Known Problems” on page 6-16
- Section 6.11, “Disconcerting Error Messages” on page 6-17
- Section 6.12, “Troubleshooting Controller Failures” on page 6-17

6.1 Controller Held in Reset, Causes, and How to Recover

This section contains the following topics:

- Section 6.1.1, “Reason Controllers Should be Failed” on page 6-2
- Section 6.1.2, “Failing a Controller in Dual/Active Mode” on page 6-3
- Section 6.1.3, “Replacing a Failed Controller” on page 6-4

The A3x00/A3500FC controllers do not detect controller failure and fail themselves. The host system via the A3x00/A3500FC drivers or the user must make the decision to fail a controller. Failing controllers is only possible in a system with redundant controllers.

The redundant array controller architecture was developed on the premise that the host system is best able to determine when a subsystem component has failed. A controller is failed if it is held in a hardware reset state. A user should fail a controller if there is cause for concern with regard to the controller’s hardware. If the controller is failed (for example held in a hardware reset state) it will not be able to access any data on the disk drives.

6.1.1 Reason Controllers Should be Failed

- Unresponsive controller

An array controller may become unresponsive as a result of a controller host chip failure, a loss of power to one of the controllers, or a controller hardware failure. The controller should always be reset, and be given adequate time cycle to through its reset logic before taking any further action.

An unresponsive controller’s typical symptoms include selection time-outs and/or continuous command time-outs. The host should first attempt to revive the controller from a possible hung state via a bus reset. If this fails, the host should continue to access the configured LUN via the alternate controller, and fail this controller.

- Obtrusive Controller

An obtrusive array controller is one which interferes with the normal operation of its alternate. This may be the result of a failing data path component of one of the array controllers, an array controller drive side SCSI bus failure, or a failing disk drive that has not been marked failed yet.

The symptoms of obtrusive array controllers may be the successful completion of some commands, particularly non data access commands. Many data access commands may fail on one or both controller in the subsystem. Other symptoms include frequent command time-outs amidst many successful command operations.

- **Failed Inter-controller Communication (ICON) Path**

Redundant controllers rely on the ICON channel, which may be a dedicated Application-specific Integrated Circuit (ASIC). The failed inter-controller communication path condition occurs when the communication path between the two array controllers has failed due to a communication line failure on one of the controllers. The controllers have the communications capability to function in the active/passive mode without the use of the ICON channel.

Diagnostics are run on the ICON channel at power up. Therefore, ICON channel failures maybe detected at Start Of Day (SOD). If you are going through the serial port and such a failure was detected, the controller's Diagnostic Manager will provide the option to the user to abort the power up sequence and replace the controller.

The failure of the ICON channel may cause some of the following situations to arise:

- Switching from any mode of operation to dual/active mode is not allowed.
- Logical unit ownership transfers are not allowed.
- Changes in drive status for example from Optimal to Failed.
- Changes in LUN status for example from Degraded to Optimal.
- The addition of a logical unit to one controller will not be seen by the other controller until the next reset or power up causes both controllers to read the array configuration stored on the disk drives.
- The deletion of a logical unit owned by one controller will not be seen by the alternate controller until the next reset or power up.

6.1.2 Failing a Controller in Dual/Active Mode

If the host determines that a controller operating in dual/active mode has failed, it may hold this controller in reset. To prevent loss of access to the LUNs owned by this controller, the host must switch mode of operation to active/failed (passive) mode.

The host should issue a mode select to the redundant controller page to the non-failed controller. This requires setting the alternate RDAC mode parameters to failed alternate controller (0x0C).

Upon receiving the mode select, the controller will:

- Attempt to quiesce itself and its pair.
- New commands to either of the controllers will terminate with a check condition indicating that quiescence is in progress.
- Write the new controller information to DacStore.
- Hold the alternate controller in reset.
- Reset the drive buses.
- Reconfigure to become the active controller in active/passive mode.
- Return status to the host for the mode select command.

The alternate controller is held in a hardware reset state, and is inaccessible to and from the host.

6.1.3 Replacing a Failed Controller

Note – A controller that “owns” logical units should not be hot swapped. You should either fail the controller prior to removal (preferable), or switch the controller to active/passive mode.

Refer to FIN I0709 and Section 5.2.6, “Controller Card Guidelines” on page 5-12 for more information on NVSRAM in the controllers.

A controller held in a hardware reset state for example failed, will have all of its LEDs on. A passive controller in the active/passive mode flashed a pattern of 0xEE or 0x6E/0xEE (module profile will indicate if it in passive mode).

When a failed controller is replaced, the new controller will not be automatically be made operational. It will remain in the hardware reset state (failed) until the good alternate controller is directed to release it from this state.

A failed controller is unfailed when it is released from the hardware reset state. This is done to allow the host to perform diagnostics on a previously failed controller, or to release a newly replaced controller from reset.

If you have a controller failure and you are using the Recovery Guru, it is very important to follow every single step as displayed through the popup windows during the recovery process. Failure to do so can cause further problems with your RAID module. Refer to Section 2.6.1, “Cluster Information” on page 2-10 for guidelines when replacing a controller within a cluster environment.

After you have brought your controller back online, do a module profile and make sure that the failed controller is in active mode and not in passive mode. If it is in passive mode, go through the maintenance application and bring that controller

back to active mode. Once you have done that, you might have to do some LUN rebalancing between the active controllers now. Also check to make sure that the firmware level matches what was on the active controller; you might have to do a firmware upgrade on the replaced controller.

If you are going to use the `rdacutil -u/U` to unfail the controller, and use the device argument form, you must use the controller that is in active mode and not the failed controller device name. The `-U` option will not do any checks and will just try to bring the controller back online through brute force.

It has been reported from the field that sometimes during a `boot -r` that one of the controllers might be thrown into a reset (failed) mode. Because of this, if you are adding a new A3x00 or moving them around, that `/dev/dsk` or `/dev/rdisk` will not get built for the controller held in reset mode.

The simplest way to resolve this is to do the following:

- Do a module profile on that A3x00 and get the `cAtBdCsD` for that good controller.
- Execute `rdacutil -u cAtBdCsD`.
- Once it is completed, do another module profile and make sure that the controller that was held in reset is active. If it is in passive mode go through the maintenance application and make it active.
- Execute `boot -r` and the paths for that controller that was held in reset mode will now be built.
- Once the system is back up, execute a `lad` and a `healthck -a`.

Note – At this point, you might want to do some LUN balancing now that both controllers are in active mode.

6.1.4 Additional ASC/ASCQ Codes

ASC/ASCQ 0x3f/02 is not documented in either the LSI Software Interface Specifications on the engineering web page or in the online file `/etc/raid/raidcode.txt`. These codes are reported by the backend disk drives in the array. These ASC/ASCQ codes can appear on the controller console, as viewed through the serial ports. The message might look like:

```
interrupt: NOTE: interrupt: 00000000 29 00 00 00 0000interrupt:
interrupt: WARN: interrupt: Sense data from device 00100001:
SKEY/ASC/ASCQ 06/3f/02 interrupt
```

The additional ASC/ASCQ codes are the following:

- 3f/00 target operating conditions have changed
- 3f/01 microcode has been changed
- 3f/02 changed operating definition
- 3f/03 inquiry data has changed

6.2 LUNs Not Seen

There are many possible causes, but the usual scenario is after reconfiguration of the system:

- After upgrade of RAID Manager 6, see `bug 4118532`.
- After upgrade of Solaris: usually `sd.conf` is lost causing LUNs above 8 are no longer seen. This is described in the *Sun StorEdge RAID Manager 6.22 Release Notes*.
- With RAID Manager 6.22, upgrading to Solaris 8 requires patch no. 108553. If you were running with patch no. 108334, it should be removed first.
- Removing or adding HBA's can cause some controllers not to be seen as described in `bug 4295322`.
- The creation of a LUN on the A3500FC in a multi-host environment requires a few extra steps to ensure that the device is properly built on both hosts. When a LUN is created with RAID Manager 6, the devices for RAID Manager and Solaris are properly configured. However, RAID Manager 6 does not configure the devices on the other host. You must configure the devices on the other host manually. See `bug 4336091` for further details.
- Problems with adding greater than 8 LUN support can cause some or all the LUNs not to be seen. FIN I0589-1 describes the proper way to increase the number of LUNs, especially using the glm HBA (PCI SCSI). Note the `add16lun.sh` script delivered on the RAID Manager 6.1.1u1 CD-ROM was wrong and should not be used. See FIN I0589-1.
- Having VERITAS Volume Manager 2.x installed and DMP enabled can cause loss of access to LUNs when a failover occurs. This configuration is not supported because it puts data at risk.
- A related issue is loss of communication with the module `rdac`. `rdac` cuts off communication when it can't access a LUN when there is a good path in `/dev/osa` for it. This can happen if LUN creation is terminated prematurely, either by operator intervention (`kill -9` or `control-C`), by system panic or power failure. `rdac_disks` usually corrects this problem or a reconfiguration reboot.
- Adding 17 LUNs to a module connected by the FC-PCI HBA, which only supports 16 LUNs will cause this loss of communication too, see `bug 4304898`.

- Adding drives from another A3x00 while the system is powered down can cause loss of LUN configuration as described in bug 4133673.

6.3 Rebuilding a Missing LUN Without Reinitialization

This section covers the following topics:

- Section 6.3.1, “Setting the VKI_EDIT_OPTIONS” on page 6-7
- Section 6.3.2, “Resetting the VKI_EDIT_OPTIONS” on page 6-9
- Section 6.3.3, “Deleting a LUN With the RAID Manager GUI” on page 6-9
- Section 6.3.4, “Recreating a LUN With the RAID Manager GUI” on page 6-9
- Section 6.3.5, “Disabling the Debug Options” on page 6-10

Rebuilding a missing LUN without reinitializing it can be dangerous, because data might get lost permanently. Use the following procedure only as a last resort after all attempts of recovery have failed.

The process of recreating a missing LUN or deleting a LUN if needed requires you to use the controller serial port and the RAID Manager GUI. You might need to recreate a LUN when it is missing after a reboot of a controller.

Before you begin to recreate a LUN, make sure of the following:

- A copy of the module profile prior to the disaster is available before attempting to recreate a LUN.
- Use the procedures for recreating a LUN if you know that the data is still intact.

Note – If you are recreating LUN 0 and if LUN 0 is greater than 10MB (for example 36GB RAID 5), and the module profile shows that it is only equal to 10MB, stop. Do not proceed with this procedure. You will not be able to recover LUN 0.

6.3.1 Setting the VKI_EDIT_OPTIONS

If the system has dual controllers, set the VKI_EDIT_OPTIONS on both controllers as follows.

1. **At the RAID controller shell prompt, enter:**

```
-> VKI_EDIT_OPTIONS
```

2. To enter insert mode, type:

```
i
```

Press Return or Enter.

3. Type:

```
writeZerosFlag=1
```

Press Return or Enter twice.

4. To enable debug options, type:

```
+
```

Press Return or Enter.

5. To quit, type:

```
q
```

Press Return or Enter.

6. To commit changes, type:

```
y
```

Press Return or Enter.

7. From the shell prompt type:

```
-> writeZerosFlag=1
```

8. From the shell prompt type:

```
-> writeZerosFlag
```

If the flag was set properly the output should indicate: value = 1 and you can proceed to Section 6.3.3, “Deleting a LUN With the RAID Manager GUI” on page 6-9 or Section 6.3.4, “Recreating a LUN With the RAID Manager GUI” on page 6-9. However, if the output says anything like: “new value added to table,” something was done incorrectly within the VKI_EDIT_OPTIONS. *Do not proceed.* Re-enter the VKI_EDIT_OPTIONS and remove the statement previously entered on both controllers.

6.3.2 Resetting the VKI_EDIT_OPTIONS

1. To clear the settings, type:

2. To confirm, type:

3. Return to Section 6.3.1, “Setting the VKI_EDIT_OPTIONS” on page 6-7.

6.3.3 Deleting a LUN With the RAID Manager GUI

Delete a LUN only if you determine it is necessary for your configuration.

1. Select the RAID module containing the LUN you want to delete.
2. Select the drive group containing the LUN you want to delete in the module information area.
3. Highlight the LUN to delete and press the Delete key.

Respond to all prompts to delete the LUN.

6.3.4 Recreating a LUN With the RAID Manager GUI

Use the module profile information gathered prior to the loss of a LUN to recreate the exact LUN parameters. Pay particular attention to drive order, segment size and caching parameters.

1. Launch the RAID Manager GUI.

1. From the Configuration screen, select a module from RAID Module.

When the LUN was deleted, all the drives assigned to that LUN should have been moved to the Unassigned drive area under module information.

2. Highlight the Unassigned drive icon.

Right click the Unassigned icon and select Create LUN...

3. Select the appropriate input for the RAID Level, Number of Drives, and Number of LUNs options.

Create an exact replica from the module profile. As you select drives under Number of Drives, select the drives in the exact order as before the loss of the LUN.

4. Click Options and select the segment size and the caching parameters

When you are finished, click OK. The "Create LUN" screen appears.

5. Click Create.

Creating the LUN takes a few minutes. The Configuration screen indicates the RAID Manager is formatting the LUN and then indicates the LUN is optimal.

6. After the LUN becomes optimal, exit RAID Manager 6 and shut down the host.

7. Power cycle the controllers.

8. Bring up the Host.

The LUN should be in its original state.

9. Restart RAID Manager and verify the LUN is created.

6.3.5 Disabling the Debug Options

After you create the LUN, perform these steps from the serial port using the VKI_EDIT_OPTIONS on both controllers to disable the Debug Options.

1. At the RAID controller shell prompt, enter:

```
-> VKI_EDIT_OPTIONS
```

2. To clear the settings, type:

```
c
```

Press Return or Enter.

3. To confirm, type:

```
y
```

4. To disable the options, type:

```
-
```

5. To confirm, type:

```
y
```

6. To quit, type:

```
q
```

7. To confirm, type:

```
y
```

8. When you are back at the prompt enter:

```
-> writeZerosFlag=0 sysReboot
```

or

```
-> sysReboot
```

6.4 Dynamic Reconfiguration

This section contains the following topics:

- Section 6.4.1, “Prominent Bugs” on page 6-12
- Section 6.4.2, “Further Information” on page 6-12

6.4.1 Prominent Bugs

- Bug 4356814 - Dynamic reconfiguration fails with A3500FC, Leadville drivers, Qlogic 2202 on an E10000 The resolution of this bug demonstrates that dynamic reconfiguration works on an E10000 over a PCI bus using QLogic 2202.
- Bug 4330698 - Unable to detach (dynamic reconfiguration) system board with A3x00/A3500FC connect. Indicates a recent problem with dynamic reconfiguration under Solaris 2.6, probably related to patch levels since dynamic reconfiguration worked on Solaris 2.6 when RAID Manager 6.22 released November 99. If the dynamic reconfiguration operation includes moving the non-pageable memory kernel cage additional steps must be taken.

6.4.2 Further Information

Refer to FIN I0536-2 for further information on dynamic reconfiguration.

Also refer to the following web sites for guides on dynamic reconfiguration:

<http://sunsolve5.sun.com/sunsolve/Enterprise-dr>

http://esp.west/home/projects/ssp3_5/pubs/ngdrpubs.html

For an A3000 disk array, see the "Special Handling of Sun StorEdge A3000" section under Chapter 2 "DR Configuration Issues" in the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*, 806-2249-10 for additional information.

6.5 Controller Failover and LUN Balancing Takes Too Long

- Controller failover can appear to take a long time as described in FIN I0634-1.

The latest isp driver is needed: 105600-15 for Solaris 2.6 only.

FIN I0551 describes editing `sd.conf` so that adding the controller back in doesn't take so long and to cut down the amount of time it takes to reboot (reducing `drvconfig` time).

- There is another issue with Fibre Channel cable failures.

Bug 4338906 - Rdac takes a long time to disable a controller with fiber problem. This bug describes a problem found in cluster testing.

Bug 4344061 - RAID Manager 6 application hangs for power off RAID module, reset loop corrects. This bug describes recovery applications hanging when power is re-applied.

Both of these bugs might be due to a known problem with Vixel 1000 hubs, Sun's only hub product as of October 2000. The hub doesn't propagate the link failure back to the host, so if the path is lost on the array side of the hub, there is no notification sent to the host. Resetting the loop via software will rectify the problem.

6.6 GUI Hang

Sometimes certain RAID Manager 6 applications such as Recovery or Maintenance will stop responding, either showing an hour glass or appearing to be dead. `Arraymon` polls the devices every few seconds as well as when each application is started. An unresponsive device can cause this. Killing the application by PID and restarting RAID Manager 6 is a reasonable workaround. For A3500FC connected arrays, using `luxadm` or `sansurfer` to reset the FC loop is also effective.

Also be aware that certain processes such as LUN reconstruction can take as long as five minutes to complete.

6.7 Drive Spin Up Failure, Drive Related Problems

Sometimes drives in an A3x00 disk array will fail without any apparent reason when the host is rebooted. Refer to `bug 4253002`.

This bug is fixed in the RAID Manager 6.1.1 Update 2 with patch no. 106513-5 or RAID Manager 6.22x.

There is another problem with drives failing to spin up due to the Calico chip in some Seagate drives. See patch no. 106817-3.

6.8 Phantom Controllers Under RAID Manager 6.22

There have been issues regarding the installation and configuration of Solaris with RAID Manager 6.22 and VERITAS Volume Manager. These issues regard instances of "phantom controllers" or device nodes. This can cause problems for your installation. To avoid these issues perform the installation of your system in the following order:

1. **Solaris:**
 - a. **Install Solaris.**
 - b. **Install required patches.**
2. **RAID Manager 6.22:**
 - a. **Run `pkgadd` to install the RAID Manager 6.22 packages.**
 - b. **Edit the `rmparams` file and change the line `"System_MaxLunsPerController=8"` to the number of LUNs needed.**
 - c. **Run `/etc/raid/bin/genscsiconf`.**
 - d. **Edit the `rdac_address` file to configure how you want your LUNs distributed over the controllers as well as to define which path the system is allowed to "see". Refer to the man page `rdac_address` for further details.**
 - e. **Run `init 6`.**
3. **Configure LUNS:**
 - a. **Run `/etc/raid/bin/rm6`.**
 - b. **Upgrade firmware (if needed but perform offline).**
 - c. **Set controller active/active (if needed).**
 - d. **Create LUNs.**
4. **VERITAS Volume Manager (VxVM):**
 - a. **Run `pkgadd` to install the VxVM 3.0.2 packages.**
 - b. **Run `vxinstall` to create a `rootdg`.**
 - c. **Run `/etc/raid/bin/rdac_disk`.**
 - d. **Run `init 6`.**

NOTES:

- You can run the `add16lun` or `add32lun` script that comes with RAID Manager 6.22. It will do all the steps needed for 16 or 32 LUNs support (`rdriver.conf` gets modified).
- Another new command, `rdac_disks`, cleans up the device tree so there is no confusion between VERITAS Volume Manager device tree and the `/dev/osa` device tree. If this step is omitted you will likely find phantom controllers, and that `lad` and `format` using different paths etc.
- If you are using any version prior to VxVM 3.0.2 you must disable DMP as described in the *Sun StorEdge RAID Manager 6.22 Release Notes* or FIN I0511-2.
- It is very popular out in the field for people to run the following commands: `drvconfig`, `devlinks`, `disks`, or `add_drv`. This will result in multiple RDAC links to a single device and the names will not agree with the RAID Manager path names. This situation can be corrected by the execution of `/etc/raid/bin/rdac_disks`. You can also look at the man pages for `rdac_disks` for more information about this.

6.9 Boot Delay (Why Booting Takes So Long)

Several things can be done to reduce delays in booting, especially reconfiguration reboots and `drvconfig` calls which are done by `rdac_disks` (1M) or `add_disk` (1M). `drvconfig` (1M) is also called when controllers are brought back online.

If you using A3500FC disk arrays only, edit the `rmparams` file and remove the `sd:` from the end of the variable `Rdac_NativeScsiDrivers:sd:ssd:`.

Also, clean up the `sd.conf` file so the SCSI device discovery process doesn't have to explicitly access each potential device. This is explained in FIN I0551.

Sometimes the number of device instances is so large that the drivers are spending time looking for non-existent devices. One way to clean up these instances is to use `disks -C` or under Solaris 8, `devfsadm -C`. This should not be done if the host is connected to controllers that are failed over, temporarily removed, or if the LUNs are not properly balanced because the extra device links will be removed.

Another approach that has been taken for a couple of escalations is to remove the line `rdnexus scsi` from `/etc/driver_classes` and replace `class=scsi` by `parent=rdnexus` in `rdriver.conf` as described in RFE no. 4374861. However, this approach has not been thoroughly tested.

Under Solaris 9, a delay of five to ten minutes has been reported. See bug 4630273.

6.10 Data Corruption and Known Problems

- Fujitsu 4/9-GB disk drive firmware 2848 has a bug and should be replaced using patch no. 108873.
- Turning power off a disk tray when using a RAID Manager version prior to 6.22. Although this is not supported it can be done accidentally. See [bug 4307641](#). RAID Manager 6.22 with firmware 3.1.x has a fix for this problem.
- RAID Manager 6.0 with firmware 2.4.x doesn't properly handle internal memory failures completely. It has been EOL'ed. Upgrade to at least RAID Manager 6.1.1 upgrade 2 immediately.
- Removing a controller after a power failure, in which cached data blocks are held on the controller.
- A single point of failure is created if cache mirroring is disabled while write-cache is enabled. Don't use this combination. If the controller with a cached write block fails, then the data is lost as described in "Cache Mirroring" on page 3-6.
- If an entire A3x00 disk array is powered down and VERITAS File System (VxFS) doesn't automatically disable the file systems on the array, then those files need to be manually disabled or data corruption will occur. See [bug 4326273](#).
- There are also problems in the isp driver that can cause data corruption so isp patch levels should be maintained, see [bug 4113677](#).
- There is an E10000 software problem with caching; upgrade Solaris 2.6 to patch no. KU-20.
- There are potential data corruption issues in RAID Manager 6.0, 6.1, and 6.1.1 firmware when failover occurs. Make sure that patch no. 106513-4 at least is installed. See [bug 4293936](#).

Note – Patch no. 106513-4 is not compatible with RAID Manager 6.0/6.1. Upgrade to at least RAID Manager 6.1.1 Upgrade 2.

- An error message with ASC/Q 0c/00 may appear to indicate data loss, but doesn't in reality. See [bug 4124793](#).

6.11 Disconcerting Error Messages

During `ufsdump` operations, the following `errno 5` message may appear:

```
Apr 10 22:29:40 abc unix: WARNING: The Array driver is returning an
Errored I/O, with errno 5, on Module 1, Lun 1, sector 43261180
```

This message can be ignored if the error only occurs when `ufsdump` is running. Otherwise the error needs to be further evaluated. See bug 4234852 and bug 4289725 and the related escalations. This message is not encountered when using RAID Manager 6.22x. The root cause is `ufsdump` reading past the end of the partition or reading 192 bytes which is not a proper "block read".

Make sure that the file system being dumped is not mounted and make sure it is clean, run `fsck`.

The error message `ASC/Q 0C/00 sense key 4,(6)` indicates write failure when its only a battery warning which may cause disabling caching or mirroring. Sometimes after power up, it takes a while for the battery to recharge. See `raidcode.txt` for the list of possibilities. Also see bug 4124793.

Another harmless message is:

```
unix: WARNING: kstat rcnt == 0 when exiting rung, please check
```

Upgrading to RM 6.22.1 will reduce the occurrence of this message. This message can be safely ignored. See bug 4671354 for details.

6.12 Troubleshooting Controller Failures

Refer to the information listed under Technical Information at:

<http://webhome.sfbay/A3x00>

Reference

This chapter contains the following topics:

- Section A.1, “Scripts and man Pages” on page A-2
- Section A.2, “Template for Gathering Debug Information for CPRE/PDE” on page A-3
- Section A.3, “RAID Manager Bootability Support for PCI/SBus Systems” on page A-4
- Section A.4, “A3500/A3500FC Electrical Specifications” on page A-5
- Section A.5, “Product Names” on page A-7

A.1 Scripts and man Pages

A number of scripts are available in the Tools directory of the released CD. The README file in the Tools directory has a description of these scripts.

A sample copy of the README file is available in the following directory:

```
/net/artemas.ebay/export/releases/sonoma/rm_6.22/rm6_22_FCS \
/Tools/README
```

The following man pages provide supplementary information for RAID Manager 6.22 array management and administration.

- arraymon
- drivutil
- fwutil
- genscsiconf
- healthck
- hot_add
- lad
- logutil
- nvutil
- parityck
- perfutil
- raidcode
- raidutil
- rdac_disks
- rdacutil
- rdaemon
- rm6
- rmscript
- storutilsymconf
- rdac.7

A.2 Template for Gathering Debug Information for CPRE/PDE

The following template should be used when submitting information to engineering regarding problems encountered in the field with the A3x00/A3500FC.

- What is the current version of Solaris that is running on the host processor? Does the problem also occur on previous versions of Solaris, for example, Solaris 2.5.1, Solaris 7, Solaris 8, etc?
- Record the output of the "Save Module Profile" from RAID Manager 6 GUI. This output will be included but not limited to the following:
 - Controller bootware and firmware version, RAID Manager 6 version number
 - Product ID
 - Controller configuration - active/active, active/passive
 - LUN configuration - RAID level, number of pieces / piece assignment, size of LUN, caching parameters, controller assignment
- A copy of `/var/adm/messages` and the approximate time when the problem was observed
- A copy of `/etc/raid/rmparams`
- System configuration needed to replicate problem. A copy of the explorer data or the following:
 - A copy of the output from `prtdiag`
 - A copy of the output from `prtconf -vp`
 - A copy of `/etc/system`
 - A copy of `/kernel/drv/sd.conf`
 - A copy of `/kernel/drv/rdriver.conf`
 - A copy of `/kernel/drv/rdnexus.conf`
 - A copy of the output from `showrev -p | sort + 1 - 2`
 - A copy of the output from `pkginfo -l SUNWosafw SUNWosar SUNWosau`
- Exact detailed steps and commands necessary to prepare the system for problem replication
- Exact detailed steps needed to replicate the problem
- Description of the problem observed and approximate time necessary in order to replicate the problem
- Name, phone number, and email address of contact person who can answer LSI Logic, Inc. questions about problem setup

- The state of the components in the A3x00/A3500FC (for example are there any failed controllers or drives, have any cables been disconnected, etc)
- A copy of the output from RAID Manager 6 health check

Note – The engineer that is working on low level A3x00/A3500FC firmware may not be very familiar with low level system administration commands, details of the configuration, or how the system operates. The engineer will require detailed information to determine what the problem is. In general, if the problem can be narrowed down to an easily reproduced scenario or one that doesn't take a long time to replicate, the engineer will have a better chance of duplicating the problem in the lab.

- A copy of `rmlog.log` after it has been run through `logutil`

A.3 RAID Manager Bootability Support for PCI/SBus Systems

The following tables provide the results of bootability testing of RAID Manager on different versions of Solaris over PCI and SBus interfaces.

Note – The A3500FC on PCI and the A3500FC on SBus are not supported.

TABLE A-1 A1000 Bootability on PCI-Based Hosts

RAID Manager Version	Solaris 2.6	Solaris 2.7	Solaris 2.8 (02/2000)	Solaris 2.8 (07/2001)	Solaris 2.9
6.1.1	Pass	Pass	Not Supported	Not Supported	Not Supported
6.22x	Fail	Fail	Fail	Fail	Not Supported

TABLE A-2 A1000 Bootability on SBus-Based Hosts

RAID Manager Version	Solaris 2.6	Solaris 2.7	Solaris 2.8 (02/2000)	Solaris 2.8 (07/2001)	Solaris 2.9
6.1.1	Pass	Pass	Not Supported	Not Supported	Not Supported
6.22x	Pass	Pass	Pass	Fail	Not Supported

TABLE A-3 A3x00 Bootability on PCI-Based Hosts

RAID Manager Version	Solaris 2.6	Solaris 2.7	Solaris 2.8 (02/2000)	Solaris 2.8 (07/2001)	Solaris 2.9
6.1.1	Pass	Pass	Not Supported	Not Supported	Not Supported
6.22x	Pass	Pass	Fail	Fail	Not Supported

TABLE A-4 A3x00 Bootability on SBus-Based Hosts

RAID Manager Version	Solaris 2.6	Solaris 2.7	Solaris 2.8 (02/2000)	Solaris 2.8 (07/2001)	Solaris 2.9
6.1.1	Pass	Pass	Not Supported	Not Supported	Not Supported
6.22x	Pass	Pass	Pass	Fail	Not Supported

Note – A3500FC bootability is not supported.

A.4 A3500/A3500FC Electrical Specifications

Refer to Appendix A.1 “Initial Cold Start Surge Current Specifications” in the *Sun StorEdge A3500/A3500FC Hardware Configuration Guide*. Also refer to Appendix B.2 “Electrical Specifications” in the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

Also refer to Appendix B.2 “Electrical Specifications” in the *Sun StorEdge A3500/A3500FC Controller Module Guide*.

The following table provides power consumption information for a given array system configuration (minimum and maximum). The difference in power consumption at 30°C and 40°C is due to the cooling fans spinning at a higher speed at 40°C.

TABLE A-5 Power Consumption Specifications

Configuration	Power Consumption at 30°C (BTU/Watts)	Power Consumption at 40°C (BTU/Watts)
A3500 Lite using 9-GB or 18-GB disk drives (minimum configuration)	1602/469	1803/528
A3500 Lite using 9-GB or 18-GB disk drives (maximum configuration)	2847/834	3048/893
1x5 using 9-GB disk drives (minimum configuration)	2131/624	2634/772
1x5 using 9-GB disk drives (maximum configuration)	6021/1764	6524/1912
1x5 using 18-GB or 36-GB (1") disk drives (minimum configuration)	2206/646	2709/794
1x5 using 18-GB or 36-GB (1") disk drives (maximum configuration)	6472/1896	6975/2044
1x5 using 36-GB (1.6") disk drives (minimum configuration)	2476/725	2979/873
1x5 using 36-GB (1.6") disk drives (maximum configuration)	5845/1712	6348/1860
2x7 using 9-GB disk drives (minimum configuration)	3889/1139	4593/1345
2x7 using 9-GB disk drives (maximum configuration)	8868/2598	9572/2804
2x7 using 18-GB or 36-GB (1") disk drives (minimum configuration)	4039/1183	4744/1390
2x7 using 18-GB or 36-GB (1") disk drives (maximum configuration)	9500/2783	10204/2990
2x7 using 36-GB (1.6") disk drives (minimum configuration)	4579/1341	5283/1548
2x7 using 36-GB (1.6") disk drives (maximum configuration)	8621/2526	9324/2732
3x15 using 9-GB disk drives (minimum configuration)	6393/1873	7902/2315

TABLE A-5 Power Consumption Specifications (Continued)

Configuration	Power Consumption at 30°C (BTU/Watts)	Power Consumption at 40°C (BTU/Watts)
3x15 using 9-GB disk drives (maximum configuration)	18063/5293	19573/5735
3x15 using 18-GB or 36-GB (1") disk drives (minimum configuration)	6618/1939	8126/2381
3x15 using 18-GB or 36-GB (1") disk drives (maximum configuration)	19417/5689	20925/6131
3x15 using 36-GB (1.6") disk drives (minimum configuration)	7427/2176	8935/2618
3x15 using 36-GB (1.6") disk drives (maximum configuration)	17532/5137	19041/5579

A.5 Product Names

The following table is a matrix that lists the rack product, controller, and software name for a disk array product before and after certain dates.

TABLE A-6 Product Name Matrix

	Sun StorEdge A1000 Array	Before 04/98 RSM Disk Tray	After 04/98 RSM Disk Tray	Before 11/99 D1000 Disk Tray	After 11/99 D1000 Disk Tray	Released 11/99 Fibre Channel (Dilbert)	Release 11/99 Fibre Channel (Tabasco)
Rack Product Name		RSM 2000	SE A3000	SE A3500	SE A3500	SE A3500FC	SE A3500FC
Controller Name Tag	SE A1000	RSM 2000	SE A3000	SE A3000	SE A3500	SE A3500FC	SE A3500FC

Abbreviations:

- SE A1000 - Sun StorEdge A1000 array
- RSM 2000 - Sun RSM 2000 array
- SE A3000 - Sun StorEdge A3000 array
- SE A3500 - Sun StorEdge A3500 array
- SE A3500FC - Sun StorEdge A3500FC array

- SE A3500FCd - Sun StorEdge A3500FC with D1000 disk trays
- SE A3500FCr - Sun StorEdge A3500FC with RSM disk trays

Definitions:

- Rack Product Name—The marketing name for the rack. This name appears on the brochure or data sheet for the product.
- Sun StorEdge Controller Name Tag—The name tag is located on the face plate of the controller.

TABLE A-7 NVSRAM Product ID

Product Names	Product ID Strings
Sun StorEdge A1000 array	StorEDGE A1000
Sun StorEdge RSM 2000	RSM Array 2000
Sun StorEdge A3000 array	StorEDGE A3000
Sun StorEdge A3500 array	StorEDGE A3000
Sun StorEdge A3500FC with RSM trays	StorEdgeA3500FCr
Sun StorEdge A3500FC with Dilbert trays	StorEdgeA3500FCd

- Product ID (NVSRAM)—The product ID is set in NVSRAM. It is visible via the format command on LUN labels. It is also visible via RAID Manager 6 GUI. Select the Module Profile, select Controller Information, and look in the Product ID field.

Index

NUMERICS

98/01 ASC/ASCQ error code, 2-5

A

A3x00/A3500FC Commandments, 1-2

accessing the serial port, 1-5

ACES web site, 1-4

add_disk command, 6-15

adding

arrays, 2-6

arrays to a host with existing arrays, 2-6

arrays under VERITAS, 4-13

disk drives, 2-7

disk drives to existing arrays, 2-7

disk trays, 2-7

disk trays to existing arrays, 2-7

arrays, adding or moving, 2-6

ASC/ASCQ A0/00 error code, 2-2

ASC/Q 0C/00 sense key 4,(6) error code, 6-17

available documentation, 1-3

available tools and information, 1-3

B

backplane assembly, 5-7

battery support information label, 2-2

battery unit

checking, 2-2

replacement, 5-13

boot delay, 6-15

bootability support matrix, A-4

box sharing setup, 2-13

Break key, 1-6

bug filing hints, 1-6

C

cables

fiber-optic, 2-3

power, 2-2

SCSI, 2-3

cache

amount of, 5-13

configuration, 5-13

mirroring, 3-6

checking

battery unit, 2-2

controller module LEDs, 5-7

cluster

configurations, 2-10

information, 2-10

command

add_disk, 6-15

dip, 1-6

format, 4-4

fsck, 6-17

hot_add, 3-6

lad, 4-4

parityck, 3-11

rdac_disks, 6-15

rdacutil -U, 5-8

rdacutil -u, 5-8

storutil, 2-6

- sysReboot, 2-9, 5-16
- sysWipe, 2-9, 5-16
- tip, 1-6
- common problems, 6-1
- configuration
 - cache, 5-13
 - hardware, 2-1
 - RAID Manager, 3-1
 - RAID module, 3-3
 - reset, 5-16
 - software, 3-3, 4-1
- configurations
 - cluster, 2-10
 - multi-initiator, 2-10
 - SCSI, 2-10
- connecting power cables for new installation, 2-2
- controller
 - board LEDs, 5-9
 - card replacement, 5-12
 - failover taking too long, 6-12
 - held in reset, 6-2
 - settings, 3-9
 - switch settings, 2-4
- controller and disk tray switch settings, 2-4
- converting 1x5 to 2x7 or 3x15, 2-8
- cooling related problems, 5-14
- CPRE Group Europe web site, 1-4
- creation of a LUN, 3-5
- crossing SCSI cables, 2-4

D

- dacstor size (upgrades), 3-8
- data corruption
 - disconcerting error messages, 6-17
 - known problems, 6-16
- Debug Guide, 1-5
- debug information template, A-3
- device
 - quorum, 4-14
 - tree rearranged, 4-9
- devices
 - ghost, 4-5
- dip command, 1-6
- disconcerting error messages, 6-17
- disk drive
 - adding disk drives to existing arrays, 2-7
 - dummy drive, 5-15

- functionality, 5-3
- moving disk drives to existing arrays, 2-7
- related problems, 6-13
- replacement, 5-14
- spin up failure, 6-13
- support matrix, 2-13
- disk tray
 - common point of failure, 5-4
 - functionality, 5-4
 - replacement, 5-14
 - switch settings, 2-4
- DMP
 - enabling and disabling, 4-12
- documentation
 - obtaining, 1-3
 - web site, 1-3
- downloading
 - documentation, 1-3
 - RAID Manager, 1-4
- driver
 - Solaris kernel, 4-2
- dynamic multipathing (DMP), 4-12
- dynamic reconfiguration
 - further information, 6-12
 - prominent bugs, 6-12
 - related problems and workaround, 4-10

E

- electrical specifications, A-5
- Enterprise Services Storage ACES web site, 1-4
- error code
 - 98/01 ASC/ASCQ, 2-5
 - ASC/ASCQ A0/00, 2-2
 - ASC/Q 0C/00 sense key 4,(6), 6-17
- error messages, 6-17
- Escalation web site, 1-4
- ethernet port, 5-10
- extended LUN support, 3-5

F

- failed controller replacement, 6-4
- failing a controller in dual-active mode, 6-3
- fan failure message, 5-2
- FCOs, 1-6
- fiber-optic cables, 2-3

FINs, 1-6

firmware

- guidelines, 5-16

- information, 5-17

- upgrade steps, 5-18

format command, 4-4

FRU replacement, 5-10

fsck command, 6-17

G

GBIC support, 2-4

generating debug information, 4-3

ghost

- devices, 4-5

- LUNs, 4-5

GUI hang, 6-13

guidelines for replacing the controller card, 5-12

H

HA, 4-13

HA configuration using VERITAS software, 4-13

hardware installation and configuration, 2-1

HBA replacement, 5-10

HBA support matrix, 2-13

high availability, 4-13

hints for filing a bug, 1-6

hot_add command, 3-6

how to recover from a controller held in reset, 6-2

hub replacement, 5-12

I

independent controller setup, 2-13

information

- cluster, 2-10

- multi-initiator, 2-11

- SCSI daisy chaining, 2-11

installation

- hardware, 2-1

- new, 2-2

- RAID Manager, 3-1

- software, 3-2, 4-1

internal directory, 1-4

L

labeling

- volume, 4-5

lad command, 4-4

latest version of RAID Manager, 1-4

Linux, 1-6

local/remote switch, 2-3

long wave GBIC support, 2-4

loop ID, 2-4

LSI Logic web site, 1-4

LUN

- balancing taking too long, 6-12

- creation, 3-5

- creation process time, 3-8

- deletion and modification, 3-9

- general information, 3-5

- numbers, 3-6

LUNs

- ghost, 4-5

- maximum LUN support, 3-5

- not seen, 6-6

M

maintenance information, 5-1

man pages, A-2

maximum LUN support, 3-5

maximum server configurations, 2-12

Microsoft Windows 98/2000, 1-6

midplane replacement, 5-15

mirroring, cache, 3-6

moving

- arrays to a host with existing arrays, 2-6

- arrays under VERITAS, 4-13

- disk drives, 2-7

- disk drives to existing arrays, 2-7

- disk trays, 2-7

- disk trays to existing arrays, 2-7

multi-initiator information, 2-11

N

Network Storage web site, 1-4

new

- hardware installation, 2-2

- software installation, 4-2

NVSRAM settings, 3-9

O

obtaining

- Debug Guide, 1-5
- RAID Manager, 1-4
- serial cable, 1-5

obtrusive controller, 6-2

onboard SOC+, 2-12

OneStop Sun Storage Products web site, 1-4

P

parity check settings, 3-10

parityck command, 3-11

patch information, 5-17

PFA, 3-3

phantom

- controllers under RAID Manager 6.22, 6-14

power

- cables, 2-2
- sequencer, 2-3
- sequencer configuration for 3x15, 2-3
- sequencer configuration for new installation, 2-3
- sequencer replacement, 5-11

power consumption specifications, A-5

predictive failure analysis, 3-3

product name matrix, A-7

Q

QAST web site, 1-4

quorum device, 4-14

R

RAID

- module configuration, 3-3
- phantom controllers, 6-14
- reconstruction rate, 3-7
- use of RAID levels, 3-6

RAID Manager

- 6.0/6.1 not supported, 1-5
- bootability support, A-4
- commands, A-2
- installation and configuration, 3-1
- issues when upgrading from RAID Manager 6 to

6.11 or 6.22, 4-2

upgrade, 5-18

upgrading to 6.22, 1-5, 3-2

white paper, 1-6

rdac_disks command, 6-15

rdacutil -U command, 5-8

rdacutil -u command, 5-8

reason controllers should be failed, 6-2

reconstruction rate of RAID5, 3-7

recovering from a power supply thermal shutdown, 5-14

reference, A-1

replacing

battery unit, 5-13

controller card, 5-12

disk drives, 5-14

disk tray, 5-14

failed controller, 6-4

FRUs, 5-10

HBA, 5-10

hub, 5-12

interconnect cables, 5-11

midplanes, 5-15

power cords, 5-11

power sequencer, 5-11

reset configuration, 5-16

rmlog.log fan failure message, 5-2

S

scripts, A-2

scripts and man pages, A-2

SCSI

A3x00 SCSI Lite, 2-9

cables, 2-3

common point of failure with cables, 5-6

configurations, 2-10

crossing SCSI cables, 2-4

daisy chaining information, 2-11

disabled wide SCSI mode, 5-9

ID, 2-4

ID jumper settings, 5-7

reducing sync transfer rate, 5-9

SCSI bus length calculation, 2-4

SCSI bus maximum bus length, 2-3

SCSI ID conflict, 2-5

SCSI to FC-AL upgrade, 2-14

termination power jumpers, 5-7

- sd_max_throttle setting, 4-3
- SDS, 4-13
- second port on the SOC+ card, 2-13
- sequencer
 - power, 2-3
- serial cable, 1-5
- serial port access, 1-5
- service information, 5-1
- setting up 2x7 and 3x15, 2-8
- SNMP, 4-11
 - trap data, 4-11
- SOC+, 2-12
- software
 - configuration, 3-3, 4-1
 - guidelines, 5-16
 - information, 5-17
 - installation, 3-2, 4-1
 - new installation, 4-2
- Solaris kernel driver, 4-2
- Solaris x86, 1-6
- Solstice DiskSuite (SDS), 4-13
- Sonoma Engineering web site, 1-4
- specifications
 - electrical, A-5
 - power consumption, A-5
- Storage ACES web site, 1-4
- storutil command, 2-6
- Sun
 - Cluster information, 2-10
 - Download Center, 1-4
 - Software Shop, 1-4
 - StorEdge A3500/A3500FC electrical specifications, A-5
 - StorEdge A3x00/A3500FC Lite, 2-9
 - StorEdge D1000 tray common point of failure, 5-5
 - StorEdge RSM tray common point of failure, 5-5
- supported configurations, 2-11
- switch local/remote, 2-3
- sysReboot command, 2-9, 5-16
- sysWipe command, 2-9, 5-16

T

- template for filing a bug, 1-6
- template for gathering debug information, A-3
- thermal shutdown, recovering, 5-14
- tip command, 1-6

- tools and information, 1-3
- troubleshooting overview, 1-1
- tunable parameters and settings, 3-3

U

- unresponsive controller, 6-2
- upgrading
 - firmware, 5-18
 - RAID Manager 6 to 6.11 or 6.22, 4-2
 - SCSI to FC-AL, 2-14
 - to RAID Manager 6.22, 1-5, 3-2
- use of RAID levels, 3-6

V

- verifying
 - amount of cache, 5-13
 - backplane assembly, 5-7
 - controller boards, 5-8
 - D1000 FRUs, 5-7
 - D1000 tray functionality, 5-5
 - FRU functionality, 5-2
 - functionality of the disk drives, 5-3
 - functionality of the disk tray, 5-4
 - functionality of the power sequencer, 5-5
 - HBA, 5-8
 - paths to the A3x00/A3500FC, 5-8
 - RSM tray functionality, 5-5
- VERITAS
 - adding or moving arrays, 4-13
 - enabling and disabling Volume Manager DMP, 4-12
 - Volume Manager, 4-12

- volume
 - labeling, 4-5

W

- web sites, 1-4
 - dynamic reconfiguration information, 6-12
 - edist home page, 1-3
 - Enterprise Services FIN & FCO Program, 1-6
 - Enterprise Services Storage ACES, 1-4
 - Escalation Web Interface, 1-4
 - firmware information, 5-17

- LSI Logic, 1-4
- Network Storage, 1-4
- OneStop Sun Storage Products, 1-4
- patch information, 5-17
- PatchPro, 5-17
- QAST Group, 1-4
- RAID Manager 6.22 documentation, 5-17
- RAID Manager 6.22 upgrades, 1-5, 3-2
- RAID Manager software documentation, 4-11
- software information, 5-17
- Sonoma Engineering, 1-4
- Sun Cluster 2.1 documentation, 4-14
- Sun Cluster 2.2 documentation, 4-14
- Sun Cluster 2.2 field Q&A, 2-11
- Sun Cluster 3.0 documentation, 2-10
- Sun Cluster documentation, 3-4
- Sun Cluster download information, 3-4
- Sun Cluster engineering technical docs & download information, 2-10
- Sun Cluster home page, 2-10
- Sun Cluster support matrix, 2-11
- Sun Download Center, 1-4
- Sun Field Engineer's Handbook, 5-7
- Sun Shopware Shop, 1-4
- Sun StorEdge A1000 A3x00 installation supplement, 5-17
- SunSolve Early Notifier home page, 5-17, 6-12
- SunStorEdge A3500FC Lite solution, 2-9
- why booting takes so long, 6-15
- World Wide Name (WWN), 2-5