



Sun StorEdge™ 3000 Family Configuration Service 1.5 User's Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-3337-12
March 2004, Revision A

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

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

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

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

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

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

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

Netscape is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries.

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

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

Copyright © 2002-2004 Dot Hill Systems Corporation, 6305 El Camino Real, Carlsbad, Californie 92009, Etats-Unis. Tous droits réservés.

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

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

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

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

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

Netscape est une marque de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

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



Please
Recycle



Adobe PostScript

Contents

- Preface xv**

- 1. Introduction 1**

- 2. Before You Begin 3**
 - Installing the Software 3
 - Example Screens 3
 - Superuser Privileges 4
 - Sun StorEdge 3120 SCSI Array 4
 - Sun StorEdge 3000 Family JBODs 4

- 3. Setting Up Sun StorEdge Configuration Service 5**
 - Starting Sun StorEdge Configuration Service 6
 - ▼ To Start Sun StorEdge Configuration Service on a Solaris, Linux, HP-UX, and AIX Host 6
 - ▼ To Start Sun StorEdge Configuration Service on a Microsoft Windows Host 7
 - Main Window Before and After First Use 7
 - If Console Locks Up During Use 7
 - ▼ To Stop Sun StorEdge Configuration Service 7
 - The Menu, Toolbar, and Tabs 8
 - Menu Bar 8

Toolbar	9
Tabs	11
Physical (Hard) Drives	11
Navigating	11
▼ To Add Servers	12
▼ To Log In and Out	17
▼ To Select a Server to Manage a Controller	18
Assigning a Controller Through the Console	19
Changing a Server Assignment	20
▼ To Unassign the Managing Server	21
▼ To Manually Delete a Server Using the Terminal Window	22
▼ To Configure Agent Parameters	22
▼ To Enable JBOD Support	25
▼ To Verify Storage Configurations	25
▼ To Save the Logical Drive Configuration	27
▼ To Create Solaris Host Partitions	29
▼ To Create Windows NT Host Partitions	30
▼ To Create Windows 2000 and Windows 2003 Host Partitions	31
▼ To IBM AIX Host Logical Volumes	32
▼ To Create HP-UX Host Logical Volumes	33
What to Do Next	34

4. Full Configuration 35

Configuring Logical Drives and Logical Volumes	36
Logical Drives	36
Logical Volumes	36
Maximum Number of Supported Logical Drives, Logical Partitions, and LUN Assignments	36
▼ To Use Standard Configuration	37

- ▼ To Use Custom Configuration 40
- The New Configuration Option 42
 - Before You Use New Configuration 42
 - ▼ To Prepare for Logical Drives Larger Than 253 Gbyte 44
 - ▼ To Create and Partition a Logical Drive Using New Configuration 45
 - ▼ To Create and Partition a Logical Volume 51
 - ▼ To Clear a Configuration 53
 - ▼ To Log Out of the Configuration Level 54
- Host LUN Assignments 54
 - ▼ To Add (Map) a Host LUN 55
 - ▼ To Delete (Unmap) a Host LUN 56
- Configuration File 56
 - ▼ To Save a Configuration to a Backup File 57
 - Loading the Configuration 58
- 5. LUN Filtering (Fibre Channel Only) 59**
 - Overview 59
 - Assigning a LUN Filter 61
 - ▼ To Access the LUN Filter View 61
 - ▼ To Add HBA Devices Manually 62
 - ▼ To Remove Standard Host Mapping 64
 - ▼ To Map Logical Drives to Hosts 65
 - ▼ To Delete a LUN Filter 68
- 6. Monitoring the Array 69**
 - The Main Window 69
 - Online Help 71
 - Tree View of Product Configurations 72
 - Groups 72

The Monitoring Process	74
Auto Discovery Options	75
Viewing Detailed Device Information	76
View Group	77
View Server	78
View HBA Card	79
View Controller Configuration	80
Controllers Tab	81
Physical Drives Tab	81
Enclosure Info Tab	82
View FRU	82
View Controller Parameters	83
View Primary/Secondary Controller Configuration	83
Logical Drives Tab	86
View Logical Drive	86
View Physical Drive	87
View Enclosure	88
Battery Information	89
▼ To Verify the In-Service Date When Replacing a Battery	91
View FRU	92
Array Administration Progress	93
Agent Options Management	93
Event Log	94
Event Log File	95
▼ To Write Events to a Log File for an IBM AIX Host	96
Event Log Window	96
Severity Levels	98
Save Report	98

View Report	101
Out-of-Band Storage Management	101
▼ To Use Out-of-Band Management	103
▼ To Remove an Array From Out-of-Band Management	106
Managing Storage Through the Web	107
Web Browser Requirements	107
UNIX OS	107
Microsoft Windows OS	108
Setting Up the Array	108
▼ To Access the Console From the Web Browser	109
7. Maintaining the Array	111
Array Administration Activities	111
▼ To Check Parity	112
▼ To Schedule a Parity Check	113
Failed Drives	115
▼ To Automatically Rebuild a Drive Using a Standby Drive	115
▼ To Rebuild a Device Without a Standby Drive	116
▼ To Check the Progress of the Rebuilding Process	117
▼ To Manually Rebuild a Failed Drive	117
▼ To Restore a Logical Drive Configuration	119
▼ To Reset the Controller	122
▼ To Mute the Controller Beeper	122
▼ To Bring a Failed Controller Back Online	123
▼ To Convert a Dual Controller Array to a Single Controller Array	123
▼ To Display Performance Statistics	124

8. Updating the Configuration 127

- ▼ To Add a Logical Drive or Logical Volume From New Logical Drives 128
 - ▼ To Add a Logical Drive to a Logical Volume 130
- ▼ To Add a Logical Volume From Existing Logical Drives 132
- ▼ To Delete a Logical Drive or Logical Volume 133
 - The Logical Drive/Logical Volume Number 135
- ▼ To Create a Partition 136
 - The Logical Drive/Logical Volume Number 138
- ▼ To Delete a Partition 139
- ▼ To Expand the Capacity of a Logical Drive or Logical Volume 141
 - The Logical Drive/Logical Volume Number 142
- ▼ To Add SCSI Drives to an Existing Logical Drive 143
 - The Logical Drive/Logical Volume Number 144
- ▼ To Copy and Replace Member Drives 144
 - The Logical Drive/Logical Volume Number 146
- ▼ To Scan in New Hard Drives 146
- ▼ To Download RAID Controller Firmware 147
- ▼ To Upgrade Firmware and Boot Record 149

Downloading Firmware for Devices 150

- ▼ To Upgrade Firmware on Hard Drives 150
- ▼ To Upgrade Firmware on SAF-TE/SES Devices 152
- ▼ To Change Controller Parameters 153
- ▼ To Save Changed Values 154
 - Channel Tab 155
 - RS 232 Tab 156
 - Cache Tab 157
 - Disk Array Tab 160
 - Drive I/F Tab 161

- Host I/F Tab 163
- Redundancy Tab 164
- Network Tab 166
- ▼ To Mute the Controller Beeper 167
- ▼ To Assign or Change Standby Drives 167
- Available Servers 169
 - ▼ To Edit a Server Entry 169
- Updating the Object Data Manager on an IBM AIX Host 171
 - ▼ To Update the ODM 171
- A. RAID Basics 173**
 - RAID Terminology Overview 173
 - RAID Levels 179
 - Local and Global Spare Drives 185
- B. Monitoring JBODs 189**
 - ▼ To Enable JBOD Support 189
 - ▼ To View Component and Alarm Characteristics 191
 - ▼ To Download Firmware for Devices 191
 - ▼ To Discover a Drive 192
- C. Using the Cluster Configuration (SCSI Only) 197**
 - Planning the Cluster Configuration 197
 - Cluster Configuration Requirements 198
 - ▼ To Set Up Cluster Configuration 198
- D. Determining Host Worldwide Names (Fibre Channel Only) 203**
 - ▼ To Determine the WWN 203
 - ▼ To Determine the WWNN 205
 - ▼ To Determine the WWPN 206

E. Email and SNMP	207
How SNMP Works	208
Using Sun StorEdge Configuration Service to Send SNMP Traps	212
▼ To Send Email Messages for Each Server	212
▼ Setting Up Servers to Send Traps	215
▼ To Check the Community String for a Microsoft Windows Host	215
▼ To Specify the Trap Recipient for a Microsoft Windows Host	217
▼ To Set Up a Solaris Host	218
▼ To Set Up a Linux Host	219
▼ To Set Up an HP-UX Host	219
▼ To Set Up an IBM AIX Host	220
To Send SNMP Traps Without Using Sun StorEdge Configuration Service	221
F. Troubleshooting	223
G. Error Codes and Messages	231
Error Codes	231
Error and Status Messages	246
Installation and Program Prompts	263
Glossary	267
Index	275

Figures

FIGURE 5-1	Example of LUN Filtering	60
FIGURE 6-1	In-Band Management	102
FIGURE 6-2	Out-of-Band Management	103
FIGURE 8-1	Copying and Replacing Member Drives	144

Tables

TABLE 3-1	Main Window Toolbar Icons	9
TABLE 3-2	Main Window Tabs	11
TABLE 4-1	Maximum Number of Supported Logical Drives, Partitions, and LUN Assignments	37
TABLE 6-1	Device Status State	71
TABLE 6-2	Two-Server Group Colors	73
TABLE 6-3	Event Log Location	95
TABLE 6-4	Event Record Fields	97
TABLE 8-1	Stripe Size Per Optimization Mode	157
TABLE 8-2	Maximum Number of Disks per Logical Drive for a 2U Array	158
TABLE 8-3	Maximum Usable Capacity (Gbyte) per Logical Drive for a 2U Array	158

Preface

This guide explains how to use Sun StorEdge™ Configuration Service to configure, monitor, and manage the Sun StorEdge 3000 family arrays. For information about installing Sun StorEdge Configuration Service, refer to the *Sun StorEdge 3000 Family Software Installation Guide*.

This guide also references Sun StorEdge Diagnostic Reporter, a companion utility of Sun StorEdge Configuration Service for sending and receiving system messages from the hosts and arrays. For information about installing Sun StorEdge Diagnostic Reporter, refer to the *Sun StorEdge 3000 Family Software Installation Guide*. For information about using Sun StorEdge Diagnostic Reporter, refer to the *Sun StorEdge 3000 Family Diagnostic Reporter User's Guide*.

Unless otherwise specified, the Sun StorEdge 3120 SCSI array, Sun StorEdge 3310 SCSI array, and Sun StorEdge 3510 FC array are referred to as the *array* or *arrays*.

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

How This Book Is Organized

This book covers the following topics:

[Chapter 1](#) introduces Sun StorEdge Configuration Service features.

[Chapter 2](#) lists steps to follow to ensure that the array has been configured properly before you install and use Sun StorEdge Configuration Service.

[Chapter 3](#) contains procedures for setting up Sun StorEdge Configuration Service.

[Chapter 4](#) provides instructions for configuring the array.

[Chapter 5](#) explains how to create a LUN filter to maintain large Fibre Channel networks that share common storage (Fibre Channel only).

[Chapter 6](#) explains how to monitor the array.

[Chapter 7](#) explains how to maintain the integrity of the array.

[Chapter 8](#) explains how to change or add to the current array configuration.

[Appendix A](#) provides basic redundant array of independent disks (RAID) information.

[Appendix B](#) explains how to monitor a standalone JBOD.

[Appendix C](#) provides information about setting up a cluster configuration (SCSI Only).

[Appendix D](#) explains how to determine the host worldwide name (WWN) on a Solaris host and a Windows NT host (Fibre Channel Only).

[Appendix E](#) explains how to set up full event monitoring and email notification capabilities.

[Appendix F](#) provides troubleshooting suggestions for a list of symptoms.

[Appendix G](#) contains a list of Sun StorEdge Configuration Service error codes and messages.

The [Glossary](#) provides RAID terminology and definitions used throughout the product documentation.

Using UNIX Commands

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

- Software documentation that you received with your system
- Solaris™ operating system documentation, which is at <http://docs.sun.com>

Shell Prompts

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

Typographic Conventions

Typeface ¹	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

¹ The settings on your browser might differ from these settings.

Related Documentation

Product	Title	Part Number
Sun StorEdge 3120 SCSI array only	<i>Sun StorEdge 3120 SCSI Array Release Notes</i>	816-7955
	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual</i>	816-7956
	<i>Sun StorEdge 3000 Family Rack Installation Guide for 1U Arrays</i>	816-7964
Sun StorEdge 3310 SCSI array only	<i>Sun StorEdge 3310 SCSI Array Release Notes</i>	816-7292
	<i>Sun StorEdge 3000 Family Best Practices Manual</i>	816-7293
	<i>Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide</i>	817-3711
	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual</i>	816-7290
	<i>Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays</i>	817-3629

Product	Title	Part Number
Sun StorEdge 3510 FC array only	<i>Sun StorEdge 3510 FC Array Release Notes</i>	816-7301
	<i>Sun StorEdge 3000 Family Best Practices Manual</i>	816-7325
	<i>Sun StorEdge 3000 Family RAID Firmware 3.25 and 3.27 User's Guide</i>	817-3711
	<i>Sun StorEdge 3000 Family Installation, Operation, and Service Manual</i>	816-7300
	<i>Sun StorEdge 3000 Family Rack Installation Guide for 2U Arrays</i>	817-3629
Sun StorEdge 3120 SCSI array	<i>Sun StorEdge 3000 Family 1.5 Software Installation Guide</i>	817-3764
Sun StorEdge 3310 SCSI array	<i>Sun StorEdge 3000 Family Configuration Service 1.5 User's Guide</i>	817-3337
Sun StorEdge 3510 FC array		817-4951
	<i>Sun StorEdge 3000 Family CLI 1.5 User's Guide</i>	817-4951
	<i>Sun StorEdge 3000 Family Safety, Regulatory, and Compliance Manual</i>	816-7930
	<i>Sun StorEdge 3000 Family FRU Installation Guide</i>	816-7326

Accessing Sun Documentation

All Sun StorEdge 3000 family documentation is available online in both PDF and HTML format at the following location:

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

The following locations are specific for the SCSI array and FC array:

<http://docs.sun.com/db/coll/3120SCSIarray>

<http://docs.sun.com/db/coll/3310SCSIarray>

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

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

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

Contacting Sun Technical Support

For late-breaking news and troubleshooting tips, review the release notes for your array located in the appropriate directory:

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

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

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

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

800-USA-4SUN

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

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

508 Accessibility Features

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

Sun Welcomes Your Comments

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

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

Please include the title and part number of your document with your feedback: *Sun StorEdge 3000 Family Configuration Service 1.5 User's Guide*, part number 817-3337-12.

Introduction

This section provides a brief overview of Sun StorEdge Configuration Service features.

Note – For brevity, Sun StorEdge Configuration Service is sometimes referred to as *the program*.

Sun StorEdge Configuration Service is a sophisticated program based on the Java™ programming language, which brings together storage configuration, resource allocation, maintenance, and monitoring tools into one application for centralized administration of the array.

From a single console located on a network, system administrators can initialize network storage, change configurations, monitor status, and schedule routine maintenance through an intuitive graphical user interface (GUI).

Administrators can also dynamically allocate, reallocate, or expand capacity as storage requirements continually change with evolving network-wide storage demands.

In the event of a change in status, the program sends alerts by console display, email, or alphanumeric pager. It can also send alerts to any Simple Network Management Protocol (SNMP) enterprise manager, for example HP OpenView.

Sun StorEdge Configuration Service consists of two components:

- *Sun StorEdge Configuration Service agent* – Monitors attached disk storage devices and reports their status to the console about these devices. The agent software needs to be installed on each individual server in the network that is part of Sun StorEdge Configuration Service. The installation can be done remotely from a single workstation on the network.

- *Sun StorEdge Configuration Service console* – Contains the GUI portion of Sun StorEdge Configuration Service and displays the information reported by the agents. The console also provides remote configuration and management of the devices from a single workstation. The console needs to be installed on a computer that manages the servers on your network.

Before You Begin

This chapter contains important information you need to be aware of before you use Sun StorEdge Configuration Service. Topics covered in this chapter include:

- “Installing the Software” on page 3
- “Example Screens” on page 3
- “Superuser Privileges” on page 4
- “Sun StorEdge 3120 SCSI Array” on page 4
- “Sun StorEdge 3000 Family JBODs” on page 4

Installing the Software

For installation instructions, refer to the *Sun StorEdge 3000 Family Software Installation Guide*.

Example Screens

Many example screens are provided throughout this guide to demonstrate the program. These screens might identify either the Sun StorEdge 3310 array or the Sun StorEdge 3510 FC array in the output. Unless specifically noted that a function, and therefore its example screen, is for the Sun StorEdge 3510 FC array only, either array applies.

Superuser Privileges

You must be superuser (administrator) to run the Sun StorEdge Configuration Service console.

Sun StorEdge 3120 SCSI Array

The only Sun StorEdge Configuration Service array function supported for the Sun StorEdge 3120 array is viewing component and alarm characteristics. See [“Monitoring JBODs” on page 189](#) for procedures related to the Sun StorEdge 3120 SCSI array.

Sun StorEdge 3000 Family JBODs

The only Sun StorEdge Configuration Service array function supported for Sun StorEdge 3000 Family JBODs is viewing component and alarm characteristics. See [“Monitoring JBODs” on page 189](#) for procedures related to the Sun StorEdge 3120 SCSI array.

Note – JBOD (Just a Bunch of Disks) is an array connected directly to a server with no controllers.

Setting Up Sun StorEdge Configuration Service

This chapter describes typical procedures for setting up Sun StorEdge Configuration Service. Procedures covered in this chapter are as follows:

- [“Starting Sun StorEdge Configuration Service” on page 6](#)
 - [“To Start Sun StorEdge Configuration Service on a Solaris, Linux, HP-UX, and AIX Host” on page 6](#)
 - [“To Start Sun StorEdge Configuration Service on a Microsoft Windows Host” on page 7](#)
- [“Main Window Before and After First Use” on page 7](#)
 - [“To Add Servers” on page 12](#)
 - [“To Log In and Out” on page 17](#)
 - [“To Select a Server to Manage a Controller” on page 18](#)
 - [“To Configure Agent Parameters” on page 22](#)
 - [“To Verify Storage Configurations” on page 25](#)
 - [“To Save the Logical Drive Configuration” on page 27](#)
 - [“To Create Solaris Host Partitions” on page 29](#)
 - [“To Create Windows NT Host Partitions” on page 30](#)
 - [“To Create Windows 2000 and Windows 2003 Host Partitions” on page 31](#)
 - [“To IBM AIX Host Logical Volumes” on page 32](#)
 - [“To Create HP-UX Host Logical Volumes” on page 33](#)

Starting Sun StorEdge Configuration Service

This section explains how to start Sun StorEdge Configuration Service.

Note – You must be superuser (administrator) to run the console.

Because the console does not receive event alerts unless it is running, after configuring the array, always leave Sun StorEdge Configuration Service running in its *minimized* mode on the console workstation. Or, instead of keeping the console running, you can use Sun StorEdge *Diagnostic Reporter*, a companion utility of Sun StorEdge Configuration Service that runs as a background service that sends messages from the hosts and array to specified email addresses. For details, refer to the *Sun StorEdge 3000 Family Diagnostic Reporter User's Guide*. For another method of receiving event alerts, see [“To Send SNMP Traps Without Using Sun StorEdge Configuration Service”](#) on page 221.

Note – For Sun StorEdge Configuration Service to view and manage a SCSI or FC array, each HBA card must be connected to the primary controller.

▼ To Start Sun StorEdge Configuration Service on a Solaris, Linux, HP-UX, and AIX Host

At the command prompt, type:

```
# ssconsole
```

▼ To Start Sun StorEdge Configuration Service on a Microsoft Windows Host

For Windows 2000, choose Start → Programs → Sun StorEdge 3000 Family → Configuration Service.

For Windows 2003, choose Start → All Programs → Sun StorEdge 3000 Family → Configuration Service.

Main Window Before and After First Use

When the program initializes for the first time, the main window is blank. The Add Server window is displayed, enabling you to add servers to the Managed Servers list of the console you are using. For details on adding servers, see [“To Add Servers” on page 12](#).

Whenever you start the program *after* the first time and *after* you have selected Managed Servers, the main window displays server icons for the servers on the Managed Servers list. Click OK or Cancel to continue.

If Console Locks Up During Use

On UNIX systems, if the console locks up during use, you can stop Sun StorEdge Configuration Service, and then close and reopen the window without affecting the agent.

▼ To Stop Sun StorEdge Configuration Service

1. At the command prompt, type:


```
# ssconsole stop
```

2. Run the program again.

The window is displayed again without affecting the agent.

The Menu, Toolbar, and Tabs

It is useful to review the main window, which includes a menu bar, tabs, and a toolbar for access to key functions.



Menu Bar

The following figure shows the main menu options.

File	View	Configuration	Array Administration
Server List Setup	Event Log	Standard Configure	Rebuild
Login	View Group	Custom Configure	Parity Check
Logout	View Server	Save Configuration	Schedule Parity Check
Save Report	View HBA Card*	Load Configuration	Controller Assignment
View Report	View Controller	Configure Host/WWN (Fibre Channel only)	Controller Maintenance
Exit	View Logical Drive	LUN Filter Properties (Fibre Channel only)	Download FW for Devices
	View Physical Drive		View Peripheral Device
	View Enclosure		
	View FRU		
	Array Admin Progress		
	Agent Options Management		
	<input checked="" type="checkbox"/> Display HDD under LD		
			Help
			Contents
			About sscsConsole

*Is displayed only for out-of-band management.

Toolbar

Located below the menu bar, the toolbar provides icons that give you quick access to commonly used functions. Select an icon to activate its function. Toolbar icons can appear both as active or inactive (grayed), depending on what resources are available in the main window.

TABLE 3-1 Main Window Toolbar Icons

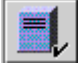





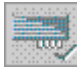
Icon	Description
	Server List Setup. Adds servers that the console manages, edits server information, or designates that an available server is managed.
	Event Log. View events such as storage device status changes, hardware status changes, or operational notifications.
	Save Event Log. Displays the Save Event Log File dialog box, which enables you to conveniently save the Event Log without having to go to the Event Log window.
	Delete Event Log. Manually deletes the contents of the eventlog.txt file. (The program accumulates events until the limit of 10,000 is reached, at which time the Event Log is automatically reduced to the most current 500 events.)

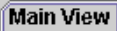

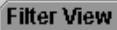
TABLE 3-1 Main Window Toolbar Icons (Continued)

Icon	Description
	Save Report. Creates a text file containing data about each of the storage components on the selected server as of the current date.
	Standard Configuration. Creates one or more logical drives with one RAID level on the selected array controller. Use when you want a predefined configuration where Sun StorEdge Configuration Service automatically configures storage.
	Custom Configuration. Provides multiple choices for configuring or reconfiguring logical drives or logical volumes with varying RAID levels on the selected array controller. Provides multiple choices for configuring or reconfiguring logical drives with varying RAID levels on the selected controller. Use when you want to manually define configuration, including setting or changing controller IDs and parameters and defining or modifying RAID sets and standby drives.

Tabs

Located below the toolbar, tabs enable you to quickly move to other Sun StorEdge Configuration Service views.

TABLE 3-2 Main Window Tabs

Tab	Description
	Click to go to the main Sun StorEdge Configuration Service window.
	Click to go to Resource Allocation View. (Reserved for future use.)
	Click to go to LUN Filter View. (Fibre Channel only.)

Physical (Hard) Drives

As the array becomes fully configured, the main window displays multiple components. The physical drives that make up the logical drives are displayed by default; however, to make the main window more manageable to navigate, you can choose not to display the physical (hard) drives by deselecting View → Display HDD under LD.

Navigating


The program follows standard Java programming language keyboard and navigation operation.

▼ To Add Servers

Before you can monitor or configure a server, you need to add it to the Managed Servers list through Server List Setup.

1. **Start Sun StorEdge Configuration Service if it is not already running.**
2. **Choose File → Server List Setup.**

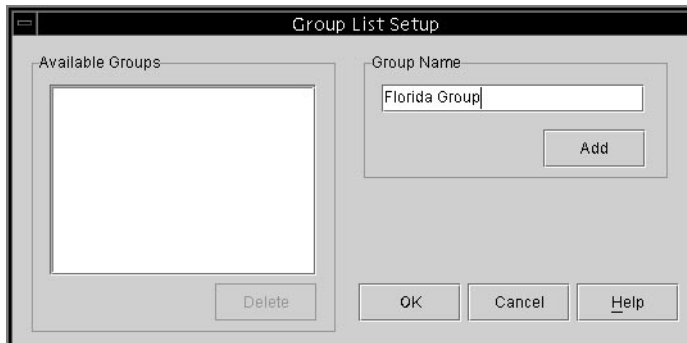
The Server List Setup window is displayed.



3. **(Optional) Organize the servers into groups:**

- a. **Click Groups in the Server List Setup window.**

The Group List Setup window is displayed.



- b. **Type a name in the Group Name field and click Add.**
- c. **To delete a group, select the group name from the Available Groups list and click Delete.**

d. When you are finished adding and deleting groups, click OK.

The Server List Setup window is displayed.

4. Add a server:

a. Click Add under the Available Servers list.

The Add Server window is displayed.

The screenshot shows the 'Add Server' dialog box with the following fields and options:

- Server name:** Lancer
- Network Information:**
 - IP Address:** 216.187.239.129
 - Socket Port:** Not assigned
 - Searching for address..
 - Get IP Addr by Name
- Monitoring Login:**
 - ssmon
 - Password: []
 - Auto discovery: Yes No
- Directory Context (For NetWare 4.x only):** []

Buttons: OK, Cancel, Help

b. Type the name of the server in the Server name field and press Return.

This is the name you use to identify the server. If this name is in your network's name server database, Sun StorEdge Configuration Service determines the server's IP address and displays it in the IP Address field.

If the program cannot find an IP address for the name, the name was either typed incorrectly or has not been recorded in the server's name database.

c. If necessary, type the server's TCP/IP address in the IP Address field.

For example, 192.168.100.1.

If the program has not already displayed the IP address (as described in the previous step), type the IP address manually.

Selecting Get IP Address by Name, located below the IP Address field, is another method of having the program search for the IP address and display it automatically. As mentioned in the previous step, it works only if your network has a name server database and you have typed the server name as it is recorded in the database. Otherwise, you must type the IP address manually.

5. Select or Deselect Automatic Discovery of servers:

- a. **Select No for Auto Discovery if you need a very secure environment where access even to server inventory data is restricted.**

When you select No, the program does not retrieve server information when starting up. The server's icon appears color-coded white (instead of active purple) to indicate it is undiscovered. When you double-click an undiscovered server, Sun StorEdge Configuration Service prompts you for the `ssmon` user password. Optionally, you can also select the server and select File → Login.

Select Yes for Auto Discovery to retrieve all information available about this server when the console is started.

- b. **If you select Yes for Auto Discovery, type the same monitor password that was typed earlier when the `ssmon` user was set up on the server (or group of servers if you have a domain or a DNS tree).**

Once the server(s) have been established using Auto Discovery, you do not need to type the `ssmon` password when you log into Sun StorEdge Configuration Service; *you automatically have monitoring privileges.*

However, whenever you choose a command to perform administration or configuration activities, you are prompted with a login dialog box to change your level of security by typing the password for either the `ssadmin` or `ssconfig` user that was established earlier.

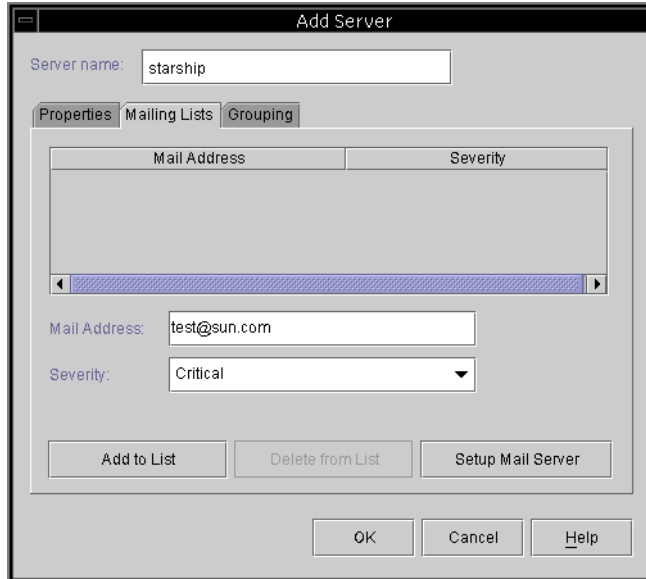
- c. **Click OK.**

6. Set email addresses:

- a. **If you want Sun StorEdge Configuration Service to send event messages using email, select the Mailing Lists tab and continue with the following directions.**

You might want to type your own email address and the addresses of selected users for the purpose of receiving information about critical events on the server.

Note – Instead of keeping the console running in the foreground, you can use Sun StorEdge Sun StorEdge *Diagnostic Reporter*, a companion utility of Sun StorEdge Configuration Service that runs as a background service that sends messages from the hosts and array to specified email addresses. For details, refer to the *Sun StorEdge 3000 Family Diagnostic Reporter User's Guide*. To ensure that Sun StorEdge Configuration Service receives email, see [“Email and SNMP” on page 207](#) for information on setting traps. For another method of receiving event alerts, see [“To Send SNMP Traps Without Using Sun StorEdge Configuration Service” on page 221](#).



- b. For each user, type an email address in the Mail Address field.**
- c. In the Severity list box, scroll through the list of severity levels and choose from the following:**

Critical – A message that requires intervention by the network administrator, such as failure of a device, power supply, or fan.

Warning – Messages that generally indicate internal program events. If you see a large number of these messages, it might mean that there is a problem with the server or the network.

Informational – Messages about the devices on the server that do not require intervention by the network administrator.

Whatever level you choose, you receive event messages for that level and any other levels at a higher severity. If you choose Informational, for example, you are notified of any form of malfunction. Conversely, if you want to be notified of only critical situations, select Critical, and you are not notified of any Informational or Warning events.


- d. Click Add to List.**

To delete a user from the list, select the mail address and click Delete from List.

e. Indicate the mail server to be used.

Note that the Setup Mail Server button toggles with Change Mail Server, depending on whether a mail server has been defined previously.

For new setups, click Setup Mail Server. A Mail Server Setup window similar to the following is displayed.



f. Type the IP address or name of the Simple Mail Transfer Protocol (SMTP) mail server that is delivering the email messages to the destination addresses specified earlier.

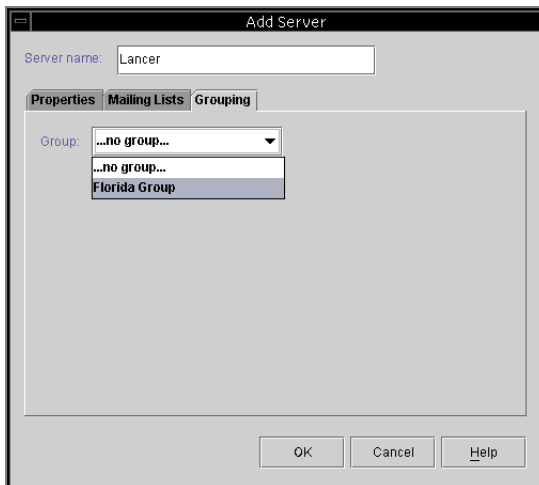
g. Click OK when you are satisfied with the information about the Mail Server window.

The Add Server window is displayed showing the Mailing Lists tab.

7. Complete the Add Server function:

a. If you want this server to be part of a group, select the Grouping tab.

The Add Server window is displayed.



b. Select the list box to view the choices available.

- c. Select the desired group and release the mouse button.
 - d. Click OK to register the change.
8. If you want to add more servers, repeat Steps 3 through 7 for each server.
 9. Move the servers you want this console to control to the Managed Servers list (see [“To Add Servers” on page 12](#)).
 - If you want this console to manage all the available servers, click Add All located at the top of the dialog box.
 - To move individual servers to the managed column, select each server individually and click Add located between the two list boxes.
 - If during the process of adding servers, modifications need to be made, see [“To Edit a Server Entry” on page 169](#).
 10. When you are finished adding servers, click OK to return to the main window.

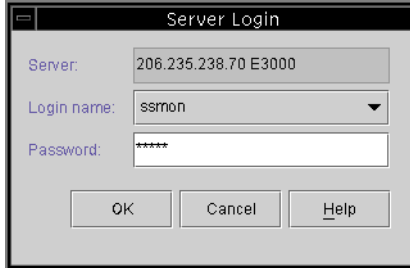
Note – If the array is attached to multiple hosts and an agent is installed on each host, each host’s IP address must be entered and added to the Managed Servers list.

▼ To Log In and Out

The log in and log out functions provide security within the program. The security levels are established by setting up three users: `ssmon`, `ssadmin`, and `ssconfig`, each with a password.

- `ssmon` represents the monitoring level of the software; it displays alerts from the controller.
- `ssadmin` represents the administration level of the software; it provides access to Rebuild, Parity Check, and Schedule Parity Check functions, as well as monitoring.
- `ssconfig` represents the configuration level; it provides access to the configuration commands and all aspects of the program.


1. To log in, choose **File** → **Login**.



Maintain continuous access only as the monitoring user. When you select a command that requires the administration or configuration level, a login dialog box is displayed and you log into that level with the appropriate password. After completing your activity, log out.

2. To log out, choose **File** → **Logout**.

When you log out from the administration or configuration level, you are given the option of logging out to the monitoring mode or logging out from the server entirely.



▼ To Select a Server to Manage a Controller

To manage and monitor a Sun StorEdge 3510 FC array, the agent needs to have access to a host logical unit number (LUN) mapped to a partition of a logical drive assigned to the primary controller of the array. The agent ignores all host LUNs mapped to a partition of a logical drive assigned to a secondary controller in a redundant configuration of an array.

Note – See [“RAID Basics” on page 173](#) for a discussion of logical drives and LUNs.

If the same array is connected to multiple servers, it is possible to have the agent running on each of these servers trying to manage and monitor the same array. Due to a restriction on monitoring commands to be sent to the array controller by only

one server at a time, some monitoring commands might fail if sent simultaneously by multiple servers. This could cause inaccurate reporting or the processes to stop responding. To prevent this from happening, the agent can be configured to enable/disable array monitoring on a server.

Assigning a Controller Through the Console

The following procedure can be performed from the console to assign server control to an array controller.



Caution – Sun StorEdge Configuration Service can monitor and manage up to 32 arrays at one time. However, console response time can decrease as the number of arrays increases.

1. **Make sure all directly attached servers were added following the directions in “To Add Servers” on page 12.**

This ensures that both Host names are present under Controller Assignments.

2. **From the main window, choose Array Administration → Controller Assignment.**

The Assign Server to Manage a RAID Controller window is displayed.

3. **Select a server from the Server to manage this controller list and click Apply.**

This enables the selected server to manage an array controller. It also disables all other servers listed from managing the same array.

A StorEdge 3000 family controller can only be managed by one host at a time. This dialog box allows you to assign the host you want to manage the controller that has a host LUN on a LD/LV assigned to the primary controller.

Controller Serial Number List				
Unique ID	Pri SN	Sec SN	Enabled Server	Disabled Server
100217	326423	326435	B52	
000198	3341258	3341275	B52	
0007eb	3357086		B52	

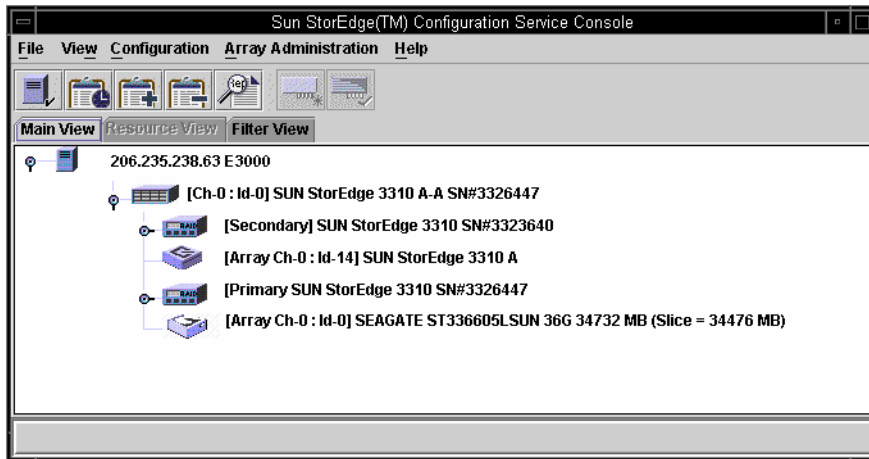
Select Server for Managing Controller

Unique ID: Controller SN:

Server to manage this controller:

4. **Click Close to close the view.**

The main window shows the controller device icon only under the server where the monitoring is enabled for this array.




Changing a Server Assignment

When you assign a server to manage a controller, information regarding the managing server is produced. The program stores this information on a controller and uses it to keep track of the managing server. In the event a server is shut down, for maintenance for example, and you try to assign another server to manage the controller, the program reads the stored server information from the controller and warns you that the controller is being managed already.

Note – You can override the current server assignment if you want to force a new server assignment.

You can view the managing server name by choosing the firmware application menu called "view and edit Host luns → Edit Host-ID/WWN Name List."



The server name is stored in hexadecimal values for the ASCII character set.


If you want to change a server assignment, for example, because you are moving an array to a new location, *before* you move the array, you need to unassign the managing server following the steps in [“To Unassign the Managing Server” on page 21](#).

If you have already moved the array, when the array starts up, you might see the warning message that the controller is already being managed by another server. You need to manually delete the server following the steps in [“To Manually Delete a Server Using the Terminal Window” on page 22](#). Refer to the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for your array for information about accessing the firmware application.

▼ To Unassign the Managing Server

1. From the main window, choose **Array Administration** → **Controller Assignment**.
2. Select the array controller for which you want to unassign a server.
3. From the **Server to manage this controller:** list box, select **none**.

Click **Close** to the following warning that is displayed confirming that the selected controller will no longer be managed by a server.



4. Select the server you want to manage the controller following the steps in [“Assigning a Controller Through the Console” on page 19](#).

▼ To Manually Delete a Server Using the Terminal Window

Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array for information about accessing the firmware application.

1. **From the Main Menu, choose “view and edit Host luns → Edit Host-ID/WWN Name List → sscsMgr → Delete Host-ID/WWN Name List.”**
2. **Choose Yes to confirm.**

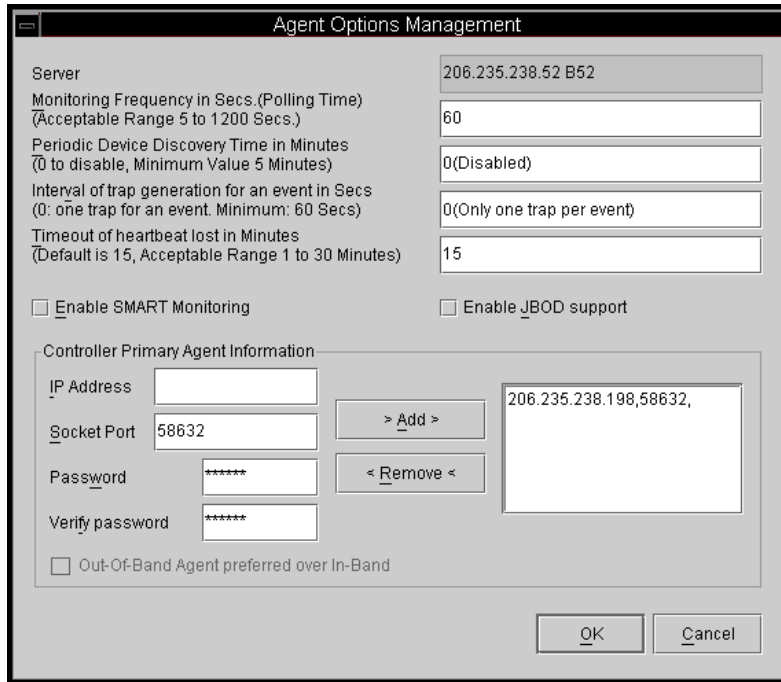
▼ To Configure Agent Parameters

This section provides steps to configure agent parameters, including polling time, periodic device discovery time, smart monitoring, out-of-band storage management, and enabling JBOD support using the View → Agent Options Management command.

1. **From the main window, choose View → Agent Options Management.**

The Agent Options Management window is displayed.

Note – Agent Options Management might not be available if one or more groups are configured and one of them is selected in the main window. To enable it, select an icon other than a group, and click View.



2. For Monitoring Frequency in Seconds, type the desired number in seconds.

This value is the interval between successive polling for any status change of devices, controllers, and enclosure monitoring by the agent. The default value is 60 seconds. If you want the console to update more frequently with device status changes, decrease this value.

Note – Increasing the polling interval could negatively impact error messaging under heavy I/O load.

3. For Periodic Device Discovery, type the value in minutes you want to check for new devices.

The periodic device discovery value is used to determine how often each device ID is scanned for a new device. The default value of 0 means do not scan for new devices. Note that device IDs are scanned less frequently as the numerical value increases.

Conversely, device IDs are scanned more frequently as the numerical value decreases. Five minutes is the minimum value.

4. **For Interval of trap generation, type the amount of time (in seconds) between the sending of each trap message.**

If the value is 60 seconds or greater, a message is sent at that interval, for that particular trap, until the event is cleared or corrected. For example, if a fan fails, a message regarding that fan failure is sent every 60 seconds until fixed.

If the value is 0, Sun StorEdge Configuration Service (and therefore Sun StorEdge Diagnostic Reporter) sends only one message regarding that particular event. For example, if a fan fails, only one email is sent.

5. **For Timeout of heartbeat lost, set the amount of time in minutes to wait between the sending of failed server messages.**

The default value is 15 minutes; the value range is 1 to 30 minutes.

6. **To Enable SMART Monitoring, select the check box.**

SMART monitoring is a method for hard drives to report predicted failures. Most disk vendors supply drives with this feature. The agent monitors this feature by issuing an unsolicited request sense. SMART monitoring can be turned off if this request causes conflicts with the underlying host operating system device drivers. For more information on SMART monitoring, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array.

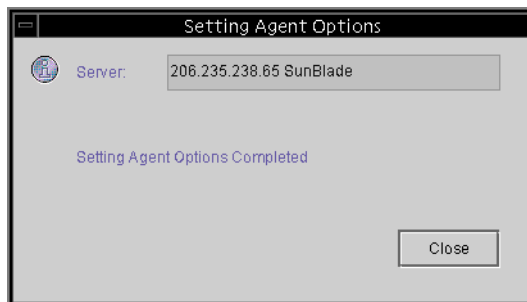
7. **To Enable JBOD support, see ["To Enable JBOD Support" on page 25](#).**

8. **The parameters under Controller Primary Agent Information, including the Password fields, pertain to out-of-band management.**

See ["To Use Out-of-Band Management" on page 103](#) for information about configuring these parameters.

9. **If you have selected any of the previous options, click OK to save your changes.**

The Setting Agent Options notification message box is displayed showing the server's TCP/IP address.



10. **Click OK to finish the procedure.**

▼ To Enable JBOD Support

Use Just a Bunch of Disks (JBOD) support only when you have a JBOD connected directly to the server. This enables you to monitor the peripheral device condition and events. If you have a JBOD connected to the RAID array, the RAID controllers monitor the JBOD condition and events for you.

Note – Enabling JBOD support could impact I/O.

1. Follow the steps in [“To Configure Agent Parameters” on page 22](#).

2. Select **Enable JBOD support**.

For details on monitoring a JBOD, see [“Monitoring JBODs” on page 189](#).

▼ To Verify Storage Configurations

Once you have installed Sun StorEdge Configuration Service and added all the servers for storage that you want managed, you need to verify the storage configurations.


Note – Most arrays are shipped preconfigured. If you want to completely remove the existing configuration and start over, see [“Full Configuration” on page 35](#). If you want to change the current configuration or add to it, see [“Updating the Configuration” on page 127](#).


1. **Make sure the server icon is online (that is, the server symbol is purple).**

During power up, Sun StorEdge Configuration Service starts automatically.

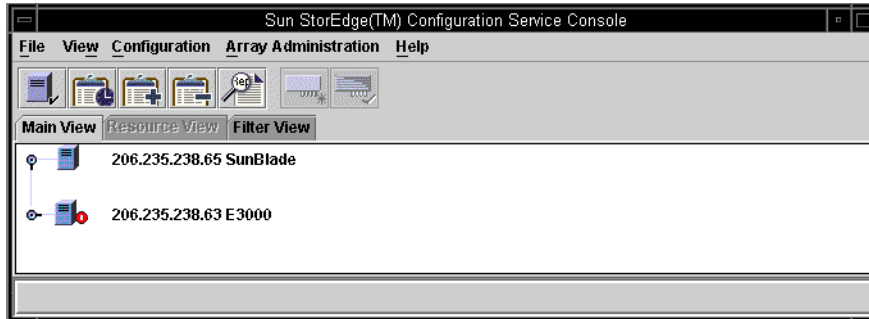
2. **Observe the main window and check the storage configuration.**


3. **If you have multiple managed servers, select the server you want to check.**

If the server icon is not purple, determine the server’s state (see [TABLE 6-1](#)). If the server’s icon has a satellite dish attached to it , the server might be in the discovery process and is available after a short delay.

4. Click the container symbol  that appears to the left of the server whose storage you want to check.

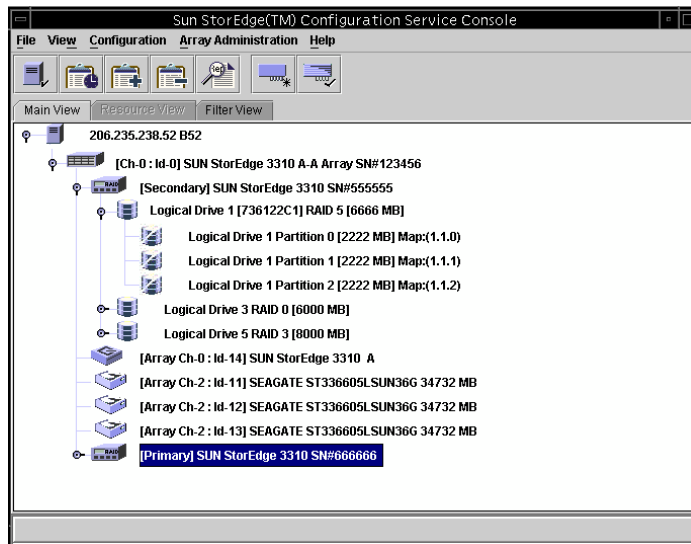
The program graphically displays each controller connected to the server as shown in the following figure.




5. To see other details, click the container symbol  next to the controller whose storage you want to check.

If the array has not yet been configured, no logical drives (LUNs for controllers) are displayed.

If the array has been fully configured, the program displays the devices and logical drives that are associated with the selected controller. It looks similar to the devices displayed in the following figure.



Selecting the container symbol  to the left of any logical drive displays its assigned physical drives.

Note – You can choose to view the physical (hard) drives that make up the logical drives by selecting or deselecting View → Display HDD under LD.

Your configuration might differ dramatically from that shown in the previous figure depending on the products you have installed.

If the array is not configured, see [“Full Configuration” on page 35](#) for instructions on configuring it.

6. **Check the RAID level and logical drive structure.**
7. **If the array is already configured and the configuration meets your needs, continue with the next section.**

If you would like to change the configuration, see [“Full Configuration” on page 35](#).

▼ To Save the Logical Drive Configuration

Even though logical drive configuration information is stored on controllers and on the physical drives attached to them, extreme circumstances, such as fire can occur, causing damage to both the controller and the drives. Keep multiple backup copies of your current logical drive (or LUN) configuration on diskette or some other form of external media other than the array. Keep at least one backup copy in a vault or secured location off site. A backup copy of the configuration enables you to restore the configuration to a new controller without having to completely reconfigure the appliance. If you do not have a backup copy of the current configuration, data could be lost.

Always save your controller configuration to a file whenever you:


- Install a new storage system enclosure or change the SCSI ID for the controller in an existing enclosure
- Replace a controller
- Reconfigure or add logical drives to a controller
- Rebuild the data from a failed drive to a standby drive

To restore a configuration from a file, see [“To Restore a Logical Drive Configuration” on page 119](#).

1. **Select the controller with the configuration you want to save.**

2. Choose Configuration → Save Configuration.

The Save Configuration window is displayed.




3. Navigate to the required drive and folder to locate the configuration file(s) to be updated, indicated by a .cfg extension.

Save the file(s) to a diskette or a drive external to the array. That is, maintain copies of these configuration files off site.

4. Specify the configuration file name and click Save.

The Save Configuration window is displayed.



5. Type a description of the configuration you are saving and click OK.

The controller configuration information is saved to a .cfg file.

▼ To Create Solaris Host Partitions

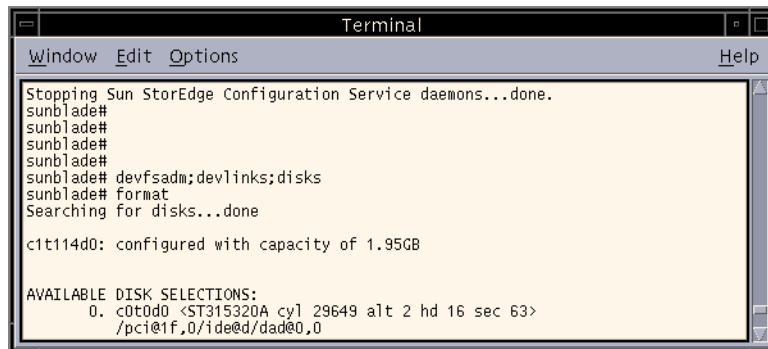
When you are satisfied with your storage configuration, you can partition the storage array for the OS.

The following steps are general guidelines. For detailed information, read about creating partitions and file systems in your Sun Solaris OS manual.

1. Label a new LUN by typing:

```
# format
```

The `format` command displays the system disk as well as other drives attached to the array.



```
Terminal
Window Edit Options Help
Stopping Sun StorEdge Configuration Service daemons...done.
sunblade#
sunblade#
sunblade#
sunblade#
sunblade# devfsadm;devlinks;disks
sunblade# format
Searching for disks...done

c1t114d0: configured with capacity of 1.95GB

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <ST315320A cyl 29649 alt 2 hd 16 sec 63>
     /pci@1f,0/ide@d/dad@0,0
```

A new LUN must be labeled with the `format` command before it can be set up for use by the array. When the `format` command is initiated, the LUNs and hard disks that are available for use are displayed.

2. Select the device to be used.

Label the device if prompted.

3. Type the word `partition` to view the existing partition table.

4. Edit the partition table as necessary and label it if changes are made.

5. Create file systems on the partitions using the `newfs` command.

6. Mount the file systems by using the `mount` command or by editing `/etc/vfstab` and using the `mountall` command.

▼ To Create Windows NT Host Partitions

Before starting this section, you need to restart the computer so that Windows NT recognizes the new storage capacity. The following steps are general guidelines; for detailed information, refer to your Windows NT documentation.

Note – This step cannot be done remotely. It must be performed at the server.

1. **Choose Start → Programs → Administrative Tools → Disk Administrator.**
2. **Select Yes in each of the Disk Administrator message boxes.**
3. **Select the logical drive by moving the cursor to the horizontal bar, and clicking the left mouse button.**

4. **Choose Disk Administrator Partition → Create.**

A prompt window is displayed, inquiring whether you want to create a partition that has a default capacity of the entire logical drive.

5. **If you want to create multiple partitions on a logical drive, enter the desired partition size.**

Repeat this step until all the logical drive capacities are used.

6. **Choose Disk Administrator Partition → Commit Changes Now.**

This function enables you to format and use the logical drive without restarting the Windows NT OS.

7. **Select the newly created logical drive partition.**

8. **Choose Disk Administrator Tools → Format. Select the NTFS (NT File System) format.**

Repeat Steps 7 and 8 to format each partition you created on the logical drive.


9. **When formatting is complete, close the Disk Administrator.**

▼ To Create Windows 2000 and Windows 2003 Host Partitions

The following steps are general guidelines; for detailed information, refer to your Windows 2000 and Windows 2003 documentation.


1. For Windows 2000, choose Start → Settings → Control Panel → Administrative Tools → Computer Management → Disk Management.

For Windows 2003, choose Start → Administrative Tools → Computer Management → Disk Management.

Make sure you can see the new drives marked by the disk icon  under Disk Management.

2. Right-click the disk for which you want to write a signature, and select Write Signature.

Right-click the disk for which you want to write a signature.



3. Select the disk for which you want to create a partition, and click OK.
4. Right-click the drive (where the drive capacity is shown), and select Create Partition.



Right-click on the drive for which you want to create a partition.

5. Respond appropriately to the partition wizard prompts.

▼ To IBM AIX Host Logical Volumes

When you are satisfied with your storage configuration, you need to create at least one logical volume on the server.

The following steps are general guidelines. For detailed information, read about creating logical volumes in your AIX OS manual.

1. **Determine that the drives are being recognized by the host by typing:**

```
# lspv
```

Ensure that the disks have been assigned a PVID (physical volume identifier). This information displays in the second column. If no PVID is assigned, the column displays "None."

2. **If no PVID is assigned, open smitty and choose Devices → Fixed Disks → Change /Show Characteristics → Assign Physical Volume Identifier.**
3. **In smitty, create a volume group.**
Choose System Storage Management → Logical Volume Manager → Volume Groups → Add a Volume Group.
4. **In smitty, create a file system.**
Choose System Storage Management → File Systems → Add/Change/Show/Delete File Systems.
5. **Mount the logical volume.**

▼ To Create HP-UX Host Logical Volumes

When you are satisfied with your storage configuration, you need to create at least one logical volume on the server.

The following steps are general guidelines. For detailed information, read about creating logical volumes in your HP-UX OS manual.

1. Determine that the drives are being recognized by the host by typing:

```
# ioscan -fnC disk
```

2. Start a System Administration Manager (sam) session.
3. Choose Disks and File Systems → Volume Groups.
4. From the Actions menu at the top of the window, click Create.
5. In the Create New Volume Group Name window, click Select New Volume Group Name, type a name for the New Volume Group, and click OK.
6. In the Create New Volume Group window, click Select Disk(s), select the drive(s) that are going to be in the Volume Group, and click OK.
7. In the Create New Volume Group window, click Define New Logical Volume(s).
 - a. In the LV name field, type a name for the logical volume.
 - b. Using the value displayed in the Approx Free Mbytes field, which specifies Mbytes left in the volume group, determine the size of the new logical volume.

Although you can create multiple logical volumes, you must create at least one. If you are creating one logical volume with the full capacity of the volume group, type the number displayed in the Approx Free Mbytes field. If you are creating multiple logical volumes, specify the size of each and type the size of the first logical volume.
 - c. In the Mount Directory field, type the directory where you want to mount the logical volume, and click Add.
 - d. To add more logical volumes, repeat Steps a-c.
 - e. When you have finished adding logical volumes, click OK.
8. In the Create New Volume Group window, click OK.
9. When you have finished creating logical volumes, close the Disk and File System window and close sam.

What to Do Next

Sun StorEdge Configuration Service (console and agent) is now installed, set up, and ready to use. See the following chapters for additional tasks:

- [“Monitoring the Array” on page 69](#), for information about how to use Sun StorEdge Configuration Service to monitor storage devices.
- [“Maintaining the Array” on page 111](#), for information about maintaining the array. It includes detailed information about parity checking, scheduling parity checks, rebuilding failed drives, and restoring a configuration from a backup file.
- [“Updating the Configuration” on page 127](#), when you want to update the configuration of the storage array. It also covers changing controller features, making or changing a standby drive, and editing the information for available servers.

Full Configuration

Sun preconfigures logical drives on the array before shipment. Read this chapter only if the array is not already configured, or if you want to *completely* remove the existing configuration and start over. If you want to make changes to the existing configuration, see [“Updating the Configuration” on page 127](#).

Full configuration includes the following topics:

- [“Configuring Logical Drives and Logical Volumes” on page 36](#)
 - [“To Use Standard Configuration” on page 37](#)
 - [“To Use Custom Configuration” on page 40](#)
 - [“To Create and Partition a Logical Volume” on page 51](#)
 - [“To Clear a Configuration” on page 53](#)
 - [“To Log Out of the Configuration Level” on page 54](#)
- [“Host LUN Assignments” on page 54](#)
 - [“To Add \(Map\) a Host LUN” on page 55](#)
 - [“To Delete \(Unmap\) a Host LUN” on page 56](#)
- [“Configuration File” on page 56](#)
 - [“To Save a Configuration to a Backup File” on page 57](#)

For the Sun StorEdge 3310 SCSI array, if you are planning to set up a cluster configuration, see [“Using the Cluster Configuration \(SCSI Only\)” on page 197](#).

The Configuration menu commands and toolbar icons might be temporarily disabled if an array administration process, such as parity checking, is running. The menu command is also deactivated when the console is refreshing its inventory on the server. A satellite dish symbol is attached to the server icon during the refresh process.



Caution – Back up any existing data to an external device before configuring the array. Reconfiguration overwrites any existing logical drive configuration. When the existing logical drive configuration is overwritten, all data is effectively erased.

Configuring Logical Drives and Logical Volumes

The information contained in this section describes how to configure logical drives and logical volumes using Standard or Custom Configuration.

Logical Drives

Depending on whether you want to use one or multiple RAID levels, logical drives can be configured using Standard or Custom Configuration.

Logical Volumes

Logical volumes are created using Custom Configuration only. However, while the ability to create and manage logical volumes remains a feature of Sun StorEdge Configuration Service for the SCSI and FC arrays, the size and performance of physical and logical drives have made the use of logical volumes obsolete. Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Avoid using them and use logical drives instead.

Maximum Number of Supported Logical Drives, Logical Partitions, and LUN Assignments

The following table lists the maximum number of logical drives, partitions per logical drive, number of partitions per logical volume, and maximum number of LUN assignments for the Sun StorEdge 3310 SCSI array and the Sun StorEdge 3510 FC array.

TABLE 4-1 Maximum Number of Supported Logical Drives, Partitions, and LUN Assignments

Product	Maximum Number of Logical Drives	Maximum Number of Partitions per Logical Drive	Maximum Number of Partitions per Logical Volume	Maximum Number of LUN Assignments
Sun StorEdge 3310 SCSI array	8	32	32	128
Sun StorEdge 3510 FC array	8	128	128	1024 (maximum number of LUNs assigned to each host is 32)

▼ To Use Standard Configuration

Use Standard Configuration Options to quickly configure all of the storage attached to an array controller using *one* RAID level. All you need to supply is the RAID controller, the optimization mode, and the RAID level. You also need to specify whether you want a standby drive. Sun StorEdge Configuration Service then automatically configures the storage for you, creating one or more logical drives, depending on the RAID level and the number of drives available. The write-back cache is selected by default. You are able to see the resulting configuration and accept or reject it before it is completed.



Caution – Back up any existing data to an external device before using Standard Configuration. Standard Configuration overwrites any existing logical drive configuration. When the existing logical drive configuration is overwritten, all data is effectively erased.



Caution – On UNIX systems, before adding new logical drives to existing configurations, you must unmount any file systems that are attached to the array.

1. **Log in as an `ssconfig` user.**
2. **Select the desired controller in the main window.**
3. **Determine whether the optimization mode of the logical drive is to be Sequential I/O (default) or Random I/O.**

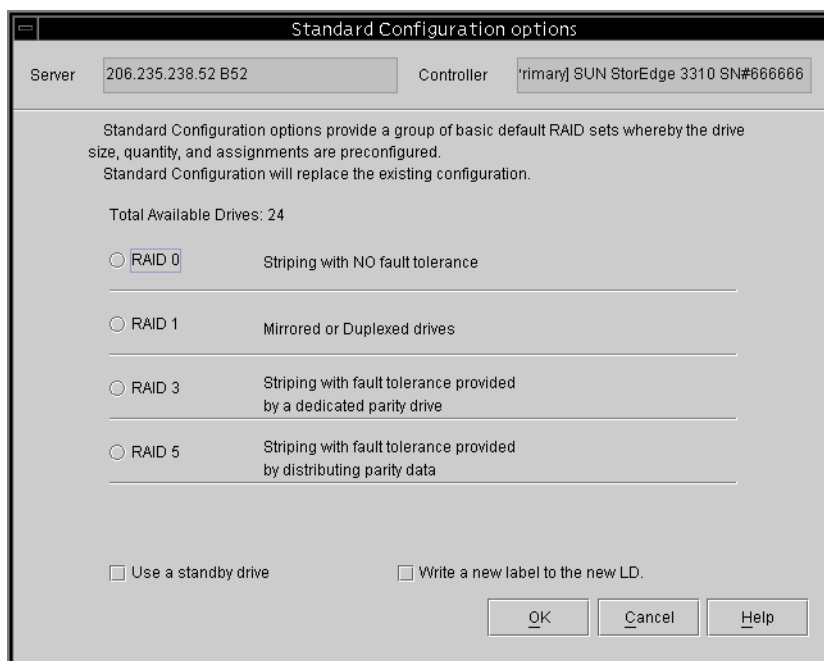
One optimization mode must be applied to all logical drives in an array. Your choice affects the maximum number of disks you can include in your array.

Optimization is set in the Cache tab of the Change Controller Parameters window. For more information on Sequential I/O and Random I/O, and for steps on how to set optimization, see [“Cache Tab” on page 157](#).

4. Choose Configuration → Standard Configure.

A Standard Configuration Warning message is displayed. Click OK to continue.

The Standard Configuration Options window is displayed. If you are not currently logged in as `ssconfig`, a login dialog is displayed.



The options on the window might be disabled, depending on the number of drives available and whether you have the Solaris OS or a Windows NT OS on the server.

The NT Cluster option creates a quorum RAID 5 logical drive (100 Mbyte) and then allocates the remaining capacity over two large RAID 5 logical drives. For this option, three or more physical drives are required.

5. Verify the server and the controller IDs displayed at the top of the window are correct.

If the server and the controller displayed at the top of the window are not correct, click Cancel to exit the configuration window and return to the main window. Select the appropriate device and select the window again.

6. Indicate whether you want to use one of the drives as a standby drive.

When you select Use a standby drive, the program deactivates the RAID levels that do not use standby drives. (For example, RAID 0 has no fault tolerance; therefore, a standby drive provides no benefit.)

7. From the RAID levels made available, select the appropriate RAID level for the logical drive you are configuring.

For definitions of the RAID levels, see “RAID Basics” on page 173.

If you select RAID 3 or 5, Standard Configuration normally configures one large logical drive. However, if the array has more than 31 physical drives, more than one logical drive is created. Also, you need a minimum of three drives to use RAID 3 or 5.

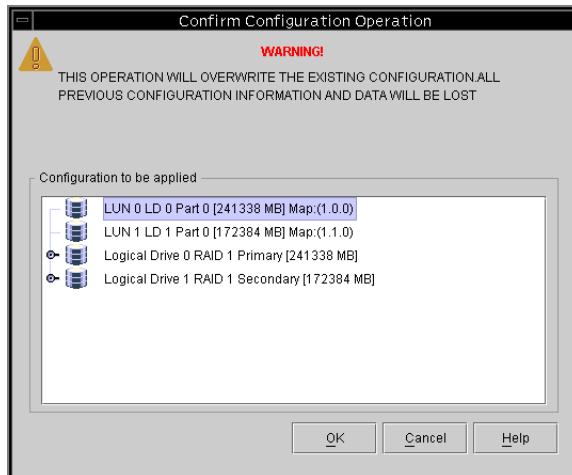
If you select RAID 1, every logical drive requires at least two physical drives. If you select four or more drives, a logical drive with a RAID 1+0 is created.

If you select RAID 0, and if the array has more than 36 drives, more than one logical drive is created.

8. (Solaris OS only). For the SCSI and FC arrays, if you want the new logical drive to be automatically labeled, which enables the OS to use the drive, click Write a new label to the new LD.

9. After you select a RAID level, click OK.

The Confirm Configuration Operation window is displayed showing the new configuration.



10. Click OK to accept the configuration as indicated; otherwise, click Cancel to return to the console.
11. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

12. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).

▼ To Use Custom Configuration

Custom Configuration gives you multiple choices for configuring or reconfiguring logical drives with *varying* RAID levels and other RAID parameters. For SCSI and FC arrays, it also gives you the option of adding newly created logical drives to logical volumes.

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).

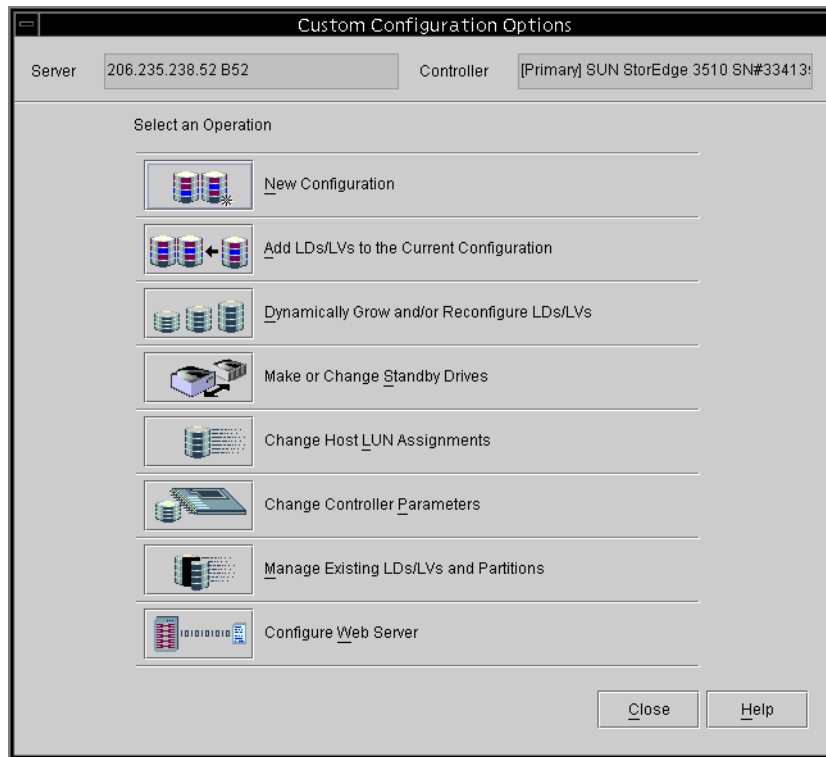


Caution – On UNIX systems, before adding new logical drives to existing configurations, you must unmount any file systems that are attached to the array.

1. Log in as an `ssconfig` user.
2. Select the desired RAID controller in the main window.

3. Choose Configuration → Custom Configure.

The Custom Configuration Options window is displayed.



- **New Configuration** – Configures new logical drives with varying RAID levels and other RAID parameters and enables you to add the new logical drives to logical volumes. The New Configuration command erases the previous configuration on an array controller.
- **Add LDs/LVs to the Current Configuration** – Adds new logical drives/logical volumes to an existing configuration. This command does not erase the other configured logical drives on the controller.
- **Dynamically Grow and/or Reconfigure LDs/LVs** – Enables you to expand the capacity of logical drives and logical volumes, add SCSI drives to existing logical drives and logical volumes, and copy and replace all member drives with drives of higher capacity.
- **Make or Change Standby Drives** – Adds one or more new standby drives to be associated with the array controller.
- **Change Host LUN Assignments** – Enables you to assign logical drives, logical volumes, and partitions to a host channel.
- **Change Controller Parameters** – Enables you to modify the controller parameters

- Manage Existing LDs/LVs and Partitions – Enables you to delete individual logical drives and logical volumes and create/delete partitions from existing logical drives and logical volumes.
- Configure Web Server (Sun StorEdge Enterprise Storage Manager Topology Reporter software required) – Enables you to configure a web server so you can manage/monitor an array through a web browser.

In this chapter, the New Configuration option is described in detail. Configure Web Server is described in [“Managing Storage Through the Web” on page 107](#). The remaining options are described in [“Updating the Configuration” on page 127](#).

The New Configuration Option

The New Configuration option enables you to customize the logical drive configuration to meet the specific needs of your environment. You can configure and partition one or more logical drives with varying RAID levels. For SCSI and FC arrays, you can then add two or more logical drives (unpartitioned) to a logical volume and divide the logical volume into a maximum of 32 partitions for SCSI and 128 partitions for Fibre Channel.

Note – On UNIX systems, if the console locks up during use, obtain the process number, and then close and reopen the window as described in [“If Console Locks Up During Use” on page 7](#).

Before You Use New Configuration

Before you use New Configuration to create a logical drive or a logical volume for the first time, it is useful to familiarize yourself with the information contained in the following steps. Knowing this information ahead of time will ease the process of creating a logical drive or volume.

1. Determine whether the optimization mode of the logical drive is to be Sequential I/O (default) or Random I/O.

One optimization mode must be applied to all logical drives in an array. Your choice affects the maximum number of disks you can include in an array.

Optimization is set in the Cache tab of the Change Controller Parameters window. For more information on Sequential I/O and Random I/O, and for steps on how to set optimization, see [“Cache Tab” on page 157](#).

- 2. Determine the total number of physical disks to be included in the logical drive, which is dependent on the optimization mode selected.**

If the security of your data is important, leave remaining disks for standby drives.

For information on the maximum number of disks and maximum usable capacity per logical drive for random and sequential optimization, see [“Maximum Number of Disks and Maximum Usable Capacity for Random and Sequential Optimization”](#) on page 158.

Note – If the logical drive is going to be larger than 253 Gbyte, see [“To Prepare for Logical Drives Larger Than 253 Gbyte”](#) on page 44.

- 3. Determine whether any standby (spare) drives are to be local or global.**

A standby drive is a drive that is marked as a spare to support automatic data rebuilding after a physical drive associated with a logical drive fails. For a standby drive to take the place of another drive, it must be at least equal in size to the failed drive. Also, the failed drive itself must be from a RAID 1, 3, or 5.

A local spare drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the local spare drive becomes a member drive and automatically starts to rebuild.

A global spare drive does not only serve one specified logical drive. When a member drive from any of the logical drives fail, the global spare drive joins that logical drive and automatically starts to rebuild. Global spares are used in the order in which they are created.

- 4. Determine the RAID level; the program automatically computes the maximum size logical drive that can be created at that RAID level.**
- 5. Determine whether you want to partition the logical drive.**

A *partition* is a logical division of the logical drive (or logical volume). Partitions enable you to create the appearance of having separate logical drives (or logical volumes) for file management, multiple users, or other purposes.

Note – Partitioning can be done during the initial configuration or after the logical drive has been created.

Note – Logical drives that have been partitioned cannot be added to a logical volume.

- 6. Understand the disk capacities displayed in the New Configuration window.**

- Max Drive Size (MB) – displays the maximum drive size per selected physical disk (total capacity of each disk).

A smaller logical drive can be created by decreasing this value. The remainder can be used later by expanding the drive (as explained in [“To Expand the Capacity of a Logical Drive or Logical Volume” on page 141](#)).

- Available Size (MB) – displays the total capacity of all disks, which is the total logical drive or logical volume size.

Note – The maximum usable capacity for a logical drive is equal to the smallest physical disk size. For example, if you add an 18-Gbyte disk and then add a 70-Gbyte disk to be part of the logical drive, the maximum usable capacity is 18 Gbyte per drive.

7. Determine whether the logical drive is to be part of a logical volume.

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).

Note – Logical drives that have been partitioned cannot be added to a logical volume.

8. Determine whether you want to partition the logical volume.

Note – Partitioning can be done during the initial configuration or after the logical volume has been created.

▼ To Prepare for Logical Drives Larger Than 253 Gbyte

The Solaris OS requires drive geometry for various operations, including `newfs`. In order for the appropriate drive geometry to be presented to the Solaris OS for logical drives larger than 253 Gbyte, you have to configure Host Cylinder/Head/Sector Mapping Configuration using the firmware application. Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array for information about accessing the firmware application.

1. In the firmware application, choose “view and edit Configuration parameters → Host-Side SCSI Parameters → Host Cylinder/Head/Sector Mapping Configuration.”
2. Choose Sector Ranges and specify 255.

3. Choose Head Ranges and specify 64.
4. Choose Cylinder Ranges and specify <65536.

▼ To Create and Partition a Logical Drive Using New Configuration

Before you create and partition a logical drive using New Configuration for the first time, it is useful to review the steps under [“Before You Use New Configuration” on page 42](#).


The following steps provide a working example of how to create a new configuration for a new logical drive. Three logical drives are selected and configured into a RAID 5 logical drive. A small logical drive is then created and partitioned.



Caution – Back up any existing data to an external device before using the New Configuration command. Sun StorEdge Configuration Service automatically initializes new logical drives defined on the array controller.

1. Log in as an `ssconfig` user.
2. Select the desired controller in the main window.
3. Choose Configuration → Custom Configure.
4. Click New Configuration.

The New Configuration Warning message box is displayed.



5. Click OK.
6. Verify that the Server and Controller displayed at the top of the window are correct.

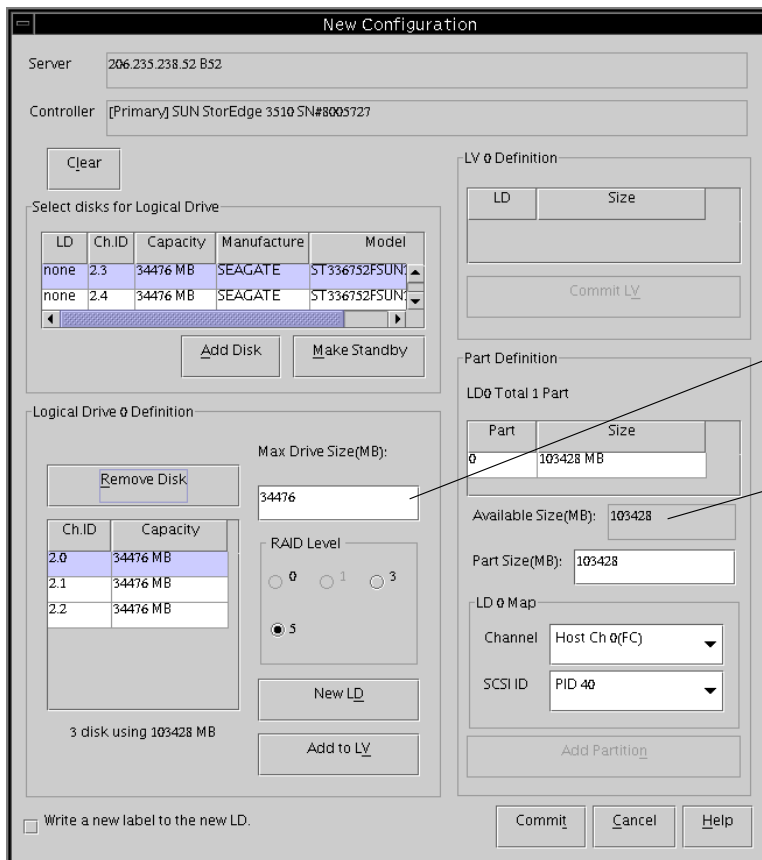
If the server and the controller displayed are not correct, click Cancel to exit the New Configuration window and return to the main window. Select the appropriate disk controller and reselect New Configuration.

7. In the Select disks for Logical Drive list box, select the first disk and click Add Disk.

Proceed to add the next two disks so that three disks are displayed in the lower list box.

If you make a mistake or change your mind, select the drive from the drives list and click Remove Disk.

Note – Because the logical drive has not been partitioned yet, the Part Size (MB) and the Available Size (MB) are equal. A single logical drive is considered to be a single partition.



Max Drive Size – displays the total capacity of each disk

Available Size – displays the total capacity of all disks, which is the total logical drive size

8. Select the RAID Level. For this example, select RAID Level 5.

This RAID level applies to all disks in this logical drive.

In this example, the new logical drive is composed of three physical disks, totaling 103428 Mbyte, as indicated in the Available Size (MB) field.

The maximum number of disks per logical drive for each RAID level for a 2U array is:

- RAID 0 - 36
- RAID 1 - 2
- RAID 1+0 - 36
- RAID 3 or 5 - 31

For RAID 1, note that if you select four or more drives, a logical drive with a RAID 1+0 is created.

9. Set the Max Drive Size.

The Max Drive Size displays the total capacity of each disk. A smaller logical drive can be created by decreasing this value.

Note – If you do not change the Max Drive Size but you do change the Partition Size, a new partition is created at the specified partition size. The remaining logical drive size capacity moves to the last partition. Remaining capacity can be used later by expanding the drive (as explained in [“To Expand the Capacity of a Logical Drive or Logical Volume” on page 141](#)). The drive capacity is no longer editable after a partition is created.

10. Select the Host channel and SCSI ID to which you would like the new logical drive to be mapped to from the Channel and SCSI ID list boxes.

Note – If you do not want to map the logical drive at this time, select Do Not Map from the Channel list box.

11. This step depends on what you want to do next:

- To create another logical drive, click New LD and repeat Steps 1-11.
- For the SCSI and FC arrays, to add this logical drive to a logical volume, click Add to LV and see [“To Create and Partition a Logical Volume” on page 51](#).
- To partition this logical drive, continue with the following steps. In this example, a small logical drive is created and partitioned.

12. To create a small logical drive, type 2000 in the Max Drive Size field.

Note – A total of 4000 Mbyte is shown for the logical drive size as opposed to 6000 Mbyte; because this is a RAID 5 level, you lose the capacity equal to 1 drive (2000 Mbyte) for parity.

The screenshot shows the 'New Configuration' window with the following sections:

- Server:** 206.235.238.52 B52
- Controller:** [Primary] SUN StorEdge 3510 SN#8005727
- Select disks for Logical Drive:** A table with columns LD, Ch.ID, Capacity, Manufacture, and Model. Two rows are visible, both with Capacity of 34476 MB and Manufacture of SEAGATE.
- Logical Drive 0 Definition:** Includes a 'Remove Disk' button, a table with columns Ch.ID and Capacity (rows 2.0, 2.1, 2.2), 'Max Drive Size(MB):' set to 2000, and 'RAID Level' set to 5.
- Part Definition:** Includes a table for 'LD0 Total 1 Part' with columns Part and Size (row 0, 4000 MB), 'Available Size(MB):' 4000, and 'Part Size(MB):' 4000.
- LD 0 Map:** Includes 'Channel' set to 'Host Ch 0(FC)' and 'SCSI ID' set to 'PID 40'.
- Buttons:** 'Commit', 'Cancel', 'Help', 'Add Disk', 'Make Standby', 'New LD', 'Add to LV', 'Add Partition', 'Commit LV'.

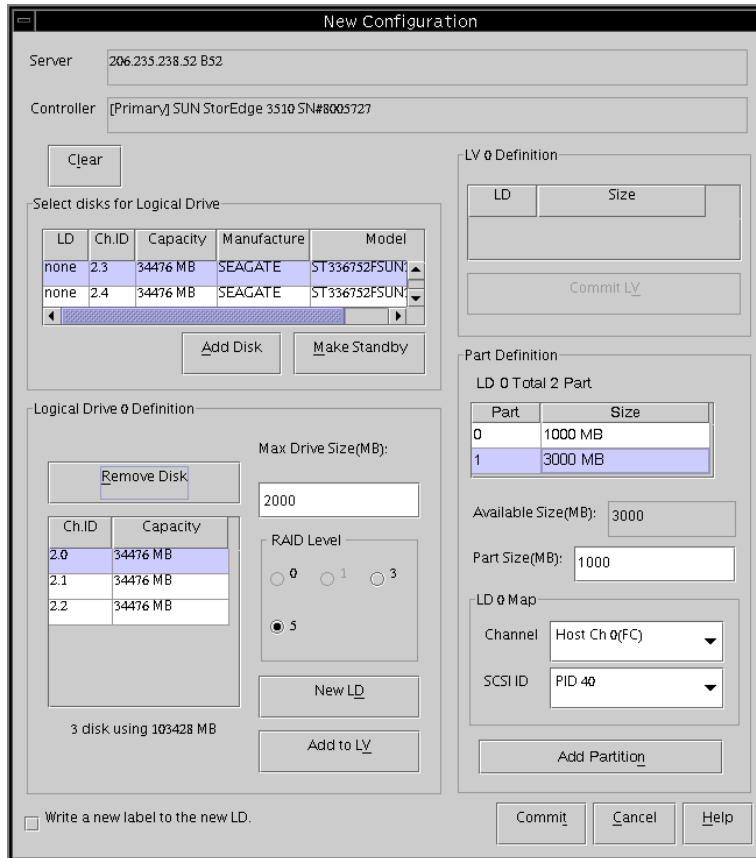
13. To create a partition, type 1000 in the Part Size field and click Add Partition.

Note – Do not partition a logical drive if you plan on adding it to a logical volume. Once a logical drive is partitioned, it cannot be added to a logical volume.

To create multiple partitions of the same size, click Add Partition as many times as partitions you want to create. You can also type the partition size in the Part Size field and multiply (*) it by the number of partitions you want to create, for example 100*128. Any remaining Mbyte is added to the last partition.

As you add partitions, the remaining capacity displayed in Available Size (MB) decreases by the amount of the partition size.

As shown in the following example, of the original 4000 Mbyte, 1000 Mbyte was allocated to Partition 0. The remaining 3000 Mbyte is then automatically moved to Partition 1. The remaining usable capacity is displayed in the Available Size (MB) field.



14. (Solaris OS only). For the SCSI and FC arrays, if you want the new logical drive to be automatically labeled, which enables the OS to use the drive, click **Write a new label to the new LD.**

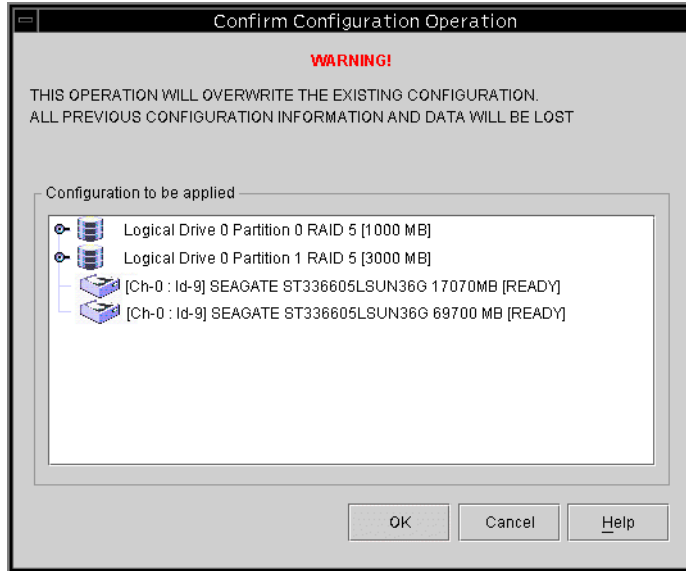
15. Click **Commit** to end the configuration of the logical drive or **New LD** to configure another logical drive.

When you click **New LD**, any available disks that haven't been used are displayed.

When the logical drive configuration for this array is complete, click **Commit**, and then click **OK**. The completed configuration is displayed.

Note – You cannot change a logical drive configuration after you click OK.

Note – During initialization, LD/LV size is displayed as 0 Mbyte.



16. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC` disk command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

17. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host”](#) on page 171.

▼ To Create and Partition a Logical Volume

A logical volume is composed of two or more logical drives and can be divided into a maximum of 32 partitions for SCSI and 128 for Fibre Channel. During operation, the host sees a nonpartitioned logical volume or a partition of a logical volume as one single physical drive.

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).

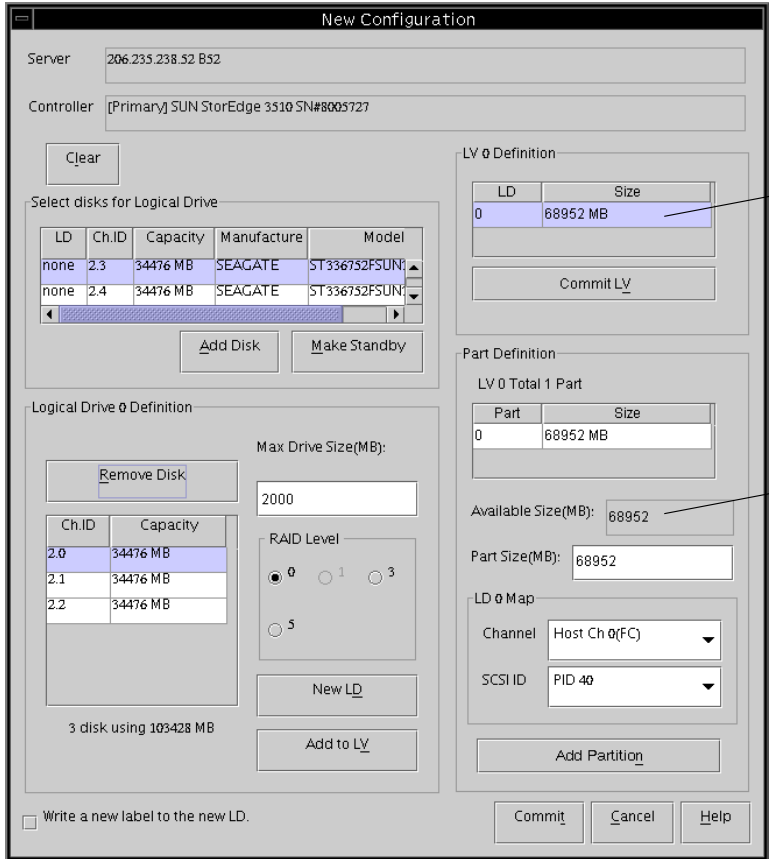
1. **Create a logical drive as described in Steps 1-11 in [“To Create and Partition a Logical Drive Using New Configuration” on page 45](#).**

Note – Do not partition the logical drive that you are adding to the logical volume. A logical drive that has been partitioned cannot be added to a logical volume.

2. **Before you click Commit, to add the logical drive to a logical volume, click Add to LV.**

The logical drive is added to the LV Definition box. The total size of the logical volume is displayed in the Available Size (MB) field.

Note – Because the logical volume has not been partitioned yet, the Part Size (MB) and the Available Size (MB) are equal. A single logical volume is considered to be a single partition.



As logical drives are added to the logical volume, they are displayed in the LV Definition box

Available Size – displays the total capacity of all disks, which is the total logical volume size

3. To create another logical drive to add to the logical volume, click **New LD**.
4. Create the logical drive and add it to the logical volume by clicking **Add to LV**. Repeat this step for every logical drive you want to add to the logical volume.
5. To create a partition, type the partition size in the **Part Size** field and click **Add Partition**.

To create multiple partitions of the same size, click **Add Partition** as many times as partitions you want to create. You can also type the partition size in the **Part Size** field and multiply (*) it by the number of partitions you want to create, for example 100*128.

As you add partitions, the remaining capacity displayed in **Available Size (MB)** decreases by the amount of the partition size.

6. **When you have finished adding logical drives to the logical volume, to create another logical volume or an individual logical drive, click Commit LV.**

When you are finished creating logical volumes and do not want to create an individual logical drive, click Commit.

Note – When you have finished creating logical volumes and want to exit the New Configuration window, if you accidentally click Commit LV instead of Commit, you will have to create another logical drive; otherwise, you have to click Cancel and configure the logical volume again.

To add or delete a logical volume to or partition a logical volume in an existing configuration, see [“Updating the Configuration” on page 127](#).

7. **(HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.**

Note – If you used system administrator manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

8. **(IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).**

▼ To Clear a Configuration

If at any time during the configuration process in the New Configuration window you are dissatisfied with the configuration, you can clear it.

1. **Click Clear to remove all physical and logical drives from the display.**
2. **Either click New LD to define a new logical drive, or click Commit.**
3. **When you click Commit, a warning prompt is displayed; click OK.**

Sun StorEdge Configuration Service initializes the entire configuration on the selected array controller.

Note – Host LUNs are automatically mapped after logical drives have completed the initialization process.

▼ To Log Out of the Configuration Level

When you have finished with configuration activities, log back in to the monitoring level of the program.

1. **Choose File → Logout.**
2. **Select Logout to monitoring mode and click OK.**

Host LUN Assignments

For the server to recognize a logical drive/logical volume, it must first be mapped to a host channel/ID. When a logical drive/logical volume is mapped to a host channel/ID, it becomes a LUN. The following section describes how to map a logical drive/logical volume to a host channel and how to remove it.

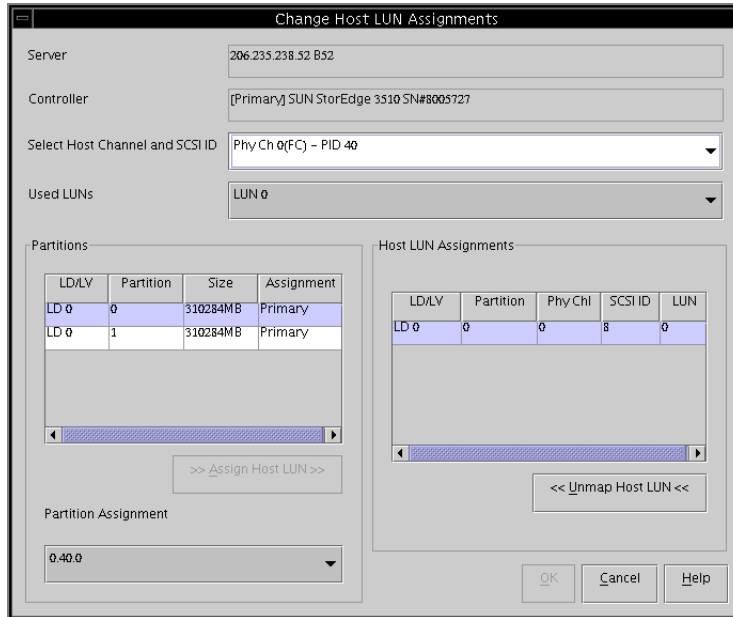
Note – When logical drives/logical volumes are first created, unless Do Not Map is selected, the host LUNs are automatically mapped after logical drives/logical volumes have completed the initialization process.

Note – On a Sun StorEdge 3310 SCSI array, there is a maximum of 128 LUN assignments. On a Sun StorEdge 3510 FC array, there is a maximum of 1024 LUN assignments (maximum number of LUNs assigned to each host ID is 32).

▼ To Add (Map) a Host LUN

1. Choose Configuration → Custom Configure.
2. Select Change Host LUN Assignments.

The Change Host LUN Assignments window is displayed.



3. Using the Select Host Channel and SCSI ID list box, select the channel and ID to which the LUN is to be assigned.

All available LUNs are displayed under Partitions. For your reference, Used LUNs displays a list of LUNs used for a specified channel and Partition Assignment displays a list of channels, IDs, and LUNs to which a partition has been assigned.

Note – When a logical drive/logical volume is added or created, it is assigned as either primary or secondary (controller assignment). A logical drive/logical volume is then initialized and mapped to a primary or secondary channel ID. If you unmap a logical drive/logical volume and remap it to another host channel ID that has a different controller assignment, *you must reset the controller*. If you unmap a logical drive/logical volume and remap it to another host channel ID that has the *same* controller assignment, you do not have to reset the controller. If a logical drive/logical volume has remaining partitions still mapped, the unmapped partitions can only be mapped to host channel IDs of the same controller assignment.

4. Select the desired partition from the LDs/LVs field, and click **Assign Host LUN** to accept the change.

To map multiple partitions (up to 32), select the first partition, scroll to the last partition and Shift-click to select everything in between. Then click **Assign Host LUN**.

As a result, **Delete Host LUN** becomes active.

5. When finished, click **OK** to save your changes and to return to the main menu.

Note – In the event that an error message is displayed when mapping a host channel, repeat Steps 1 through 5.

▼ To Delete (Unmap) a Host LUN

1. With the **Change Host LUN Assignments** window displayed, select the LUN to be deleted from the **Host LUN Assignments** field.
2. Click **Delete Host LUN**.
3. When finished, click **OK** to save your changes and to return to the main menu.

Configuration File


Keep a backup copy of the current configuration on diskette or on an external device other than the array. Even though configuration information is stored on array controllers and on the physical drives attached to them, circumstances such as a fire or a flood can occur, causing damage to both the controller and the drives. A backup copy of the configuration enables you to restore the configuration to a new controller without having to completely reconfigure the storage array.

Always save the array controller configuration to a file whenever you:


- Install a new storage system enclosure or change the SCSI ID for the controller in an existing enclosure
- Replace a controller
- Reconfigure or add logical drives to a controller
- Rebuild the data from a failed drive to a standby drive

▼ To Save a Configuration to a Backup File

1. **Select the controller with the configuration you want to save.**
2. **Choose Configuration → Save Configuration.**
The Save Configuration window is displayed.



3. **If necessary, navigate to the drive and folder where you want the configuration file to be saved.**
Save the file to a diskette or a drive external to the array.
4. **Specify the configuration file name and click Save.**
The Save Configuration window is displayed.



5. **Type a description of the configuration you are saving.**
6. **Click OK.**
The controller configuration information is saved to a `.cfg` file.

Loading the Configuration

If a drive or controller is damaged and needs to be replaced, see [“To Restore a Logical Drive Configuration” on page 119](#), which describes how to load a configuration file and restore a logical drive configuration.

LUN Filtering (Fibre Channel Only)

This chapter explains how to create a LUN filter for the Sun StorEdge 3510 Fibre Channel array to maintain large Fibre Channel networks that share common storage. Topics covered in this chapter include:

- [“Overview” on page 59](#)
- [“Assigning a LUN Filter” on page 61](#)
 - [“To Access the LUN Filter View” on page 61](#)
 - [“To Add HBA Devices Manually” on page 62](#)
 - [“To Remove Standard Host Mapping” on page 64](#)
 - [“To Map Logical Drives to Hosts” on page 65](#)
 - [“To Delete a LUN Filter” on page 68](#)

Overview

An important part of effectively maintaining a storage infrastructure depends on how accessible and secure the data is at all times. With LUN filtering support, Sun StorEdge Configuration Service provides a utility for secure centralized storage access management.

For multiple servers connected to the same FC array, LUN filtering provides an exclusive path from a server to a logical drive and essentially hides or excludes the other connected servers from seeing or accessing the same logical drive. That is, the LUN filter organizes how the FC array devices are accessed and viewed from host devices, and typically maps a FC array device to only one host so that other hosts do not access and use the same FC array device. LUN filtering also enables multiple logical drives or partitions to be mapped to the same LUN number allowing different servers to have their own LUN 0 to boot from, if needed. LUN filtering is also valuable in clarifying mapping when each HBA typically sees twice the number of logical drives when viewed through a hub.

Each Fibre Channel device is assigned a unique identifier called a worldwide name (WWN). A WWN is assigned by the IEEE and is similar to a MAC address in IP or a URL on the Internet. These WWNs stay with the device for its lifetime. LUN filtering uses this WWN to specify which server is to have exclusive use of a specific logical drive. An easy drag-and-drop method is used to map each logical drive to a host's Fibre Channel HBA card, identified with a WWN.

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

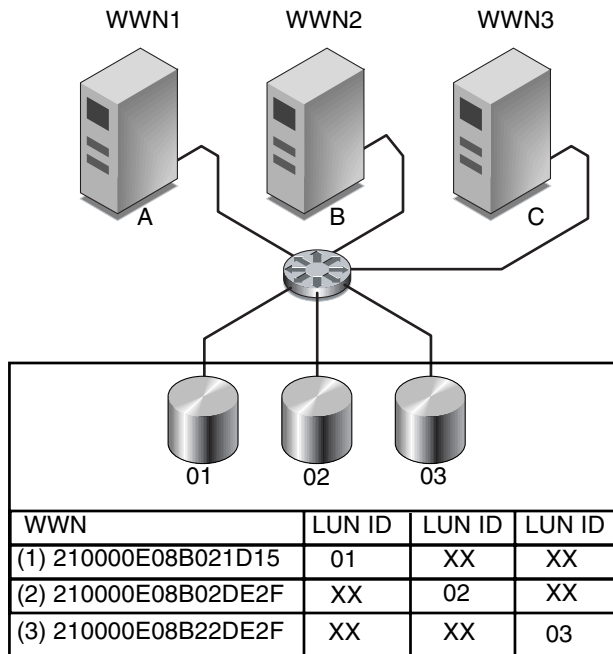


FIGURE 5-1 Example of LUN Filtering


An advantage of LUN filtering is that it allows many more hosts, like Windows NT, to attach to an array through a common Fibre Channel port and still maintain LUN security.

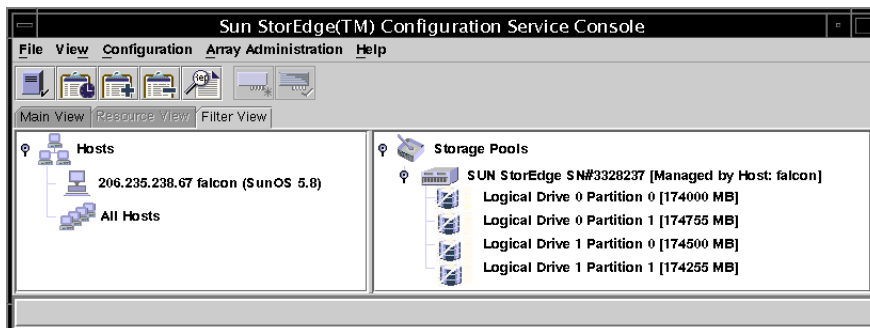
Assigning a LUN Filter

Assigning a LUN filter involves the following steps:

- Access the LUN Filter window
- Add HBA device manually (if you do not see the desired HBA)
- Remove standard mapping
- Map logical drives to hosts

▼ To Access the LUN Filter View

1. Access the LUN Filter View by clicking the Filter View tab.
2. Click the container symbol  to expand the device trees and view the servers in detail on the left and the FC array devices in detail on the right.

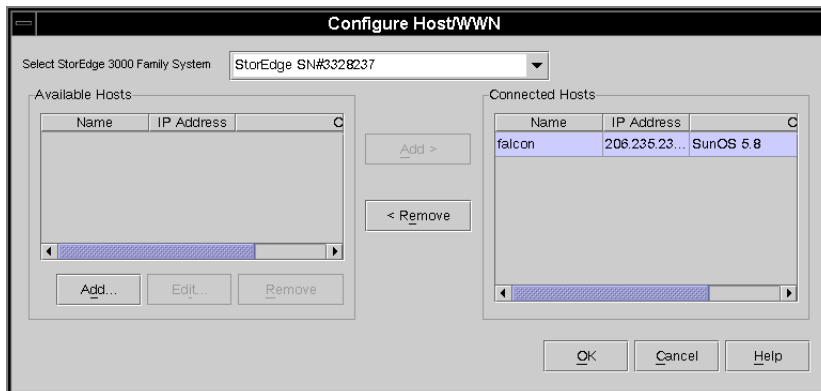


▼ To Add HBA Devices Manually


If you do not see the desired HBA device under the Hosts displayed, add it manually through the Configure Host/WWN window.

1. Choose Configuration → Configure Host/WWN.
2. In the Available Hosts list, select the host that has the HBA device to be added, and click Edit.

If the host is under the Connected Hosts list, select the host and click Remove to place the host under the Available Hosts list.




3. In the Add/Edit Host box, click Add.



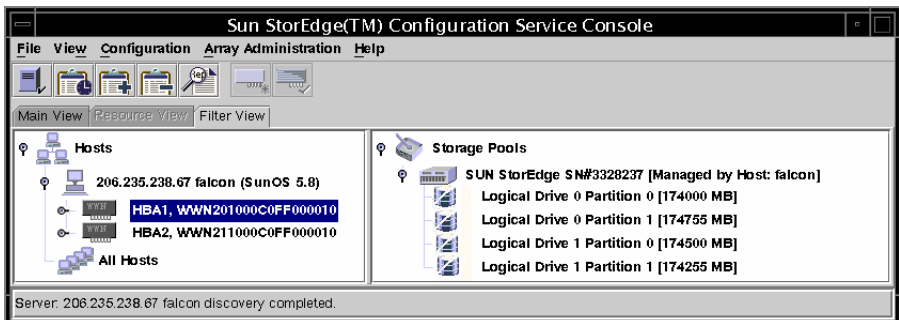
4. In the next Add/Edit HBA box, enter the new Adapter Name and the appropriate WWN, and click OK.

For details on determining the WWN, see [“Determining Host Worldwide Names \(Fibre Channel Only\)”](#) on page 203.




5. In the Configure Host/WWN window, select the edited host, and click Add to move the edited host to the Connected Hosts list, and click OK.

When you return to the LUN Filter view, the new HBA device is gray and is available for LUN filtering.



▼ To Remove Standard Host Mapping

Standard mapping allows all hosts to see all standard mapped logical drives. LUN filtering restricts the mapping to specific host(s). Before you can assign a LUN filter, you have to remove the standard host mapping (indicated by M labels ) from the FC array.

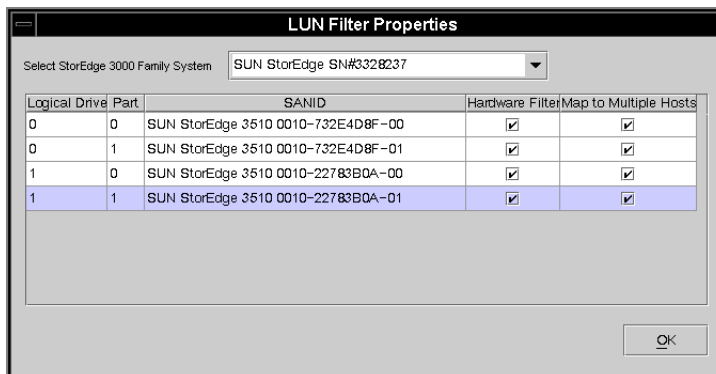
1. Click each **All Hosts** item, and drag and drop it to **Storage Pools**.
2. Click **OK** to the following message:

Operation will delete all standard maps of Logical Drive X Partition X. Do you want to continue?

A message confirming the successful deletion of the standard host LUN mapping is displayed.

3. If the host is not the primary agent for an array and you want to map the device to that host, select the device, and choose **Configuration** → **Configure Host/WWN**.
4. Select the desired host from the **Available Hosts** list, and click **Add** to add the host to the **Connected Hosts** list.
A message is displayed to confirm the addition of the host server.
5. Click **Close** to save the configuration information and exit the window.
6. If you want to map multiple hosts to an array, choose **Configuration** → **LUN Filter Properties**, and then select the array from the pulldown menu.
7. Select the **Hardware Filter** and **Map to Multiple Hosts** check boxes, and click **OK**.

Note – If you deselect the **Hardware Filter** check box, you cannot create a filter mapping; you can only create a standard map.



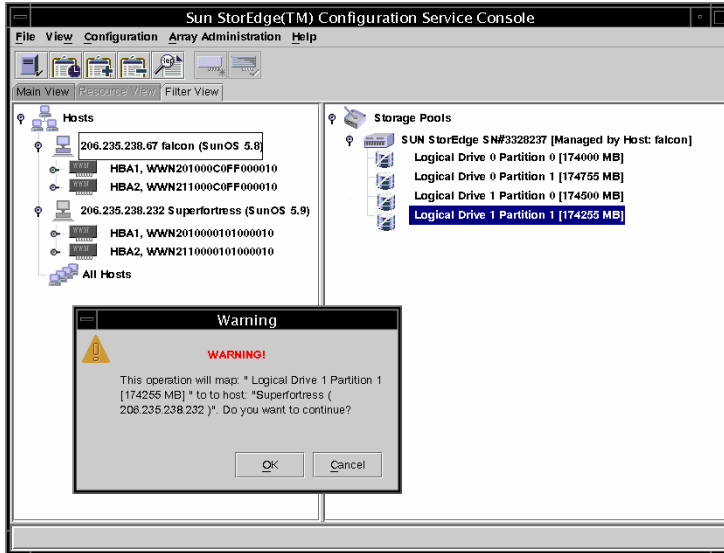
▼ To Map Logical Drives to Hosts

1. Make sure you are still in the LUN Filter view; if not, click the Filter View tab.
2. Select a logical drive from Storage Pools, and drag and drop it to the appropriate HBA card under Hosts.

If the HBA card connected to the array is not known, drag and drop the logical drive to the host.

The program prompts you to filter the array for each HBA card in the host (which means that it is visible to the HBA card attached to it and does not appear for the other HBA cards).

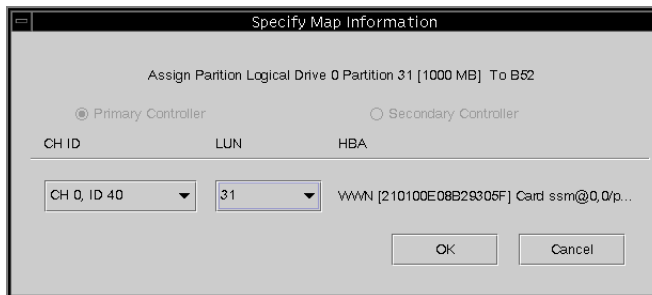
- When the prompt is displayed to confirm the mapping, click OK.



The Specify Map Information window is displayed. Note that the program automatically maps the logical drive to the first HBA card listed in the host if there is more than one HBA card.

- Assign the logical drive to a primary or secondary controller and with the desired channels and LUN ID numbers, and then click OK.

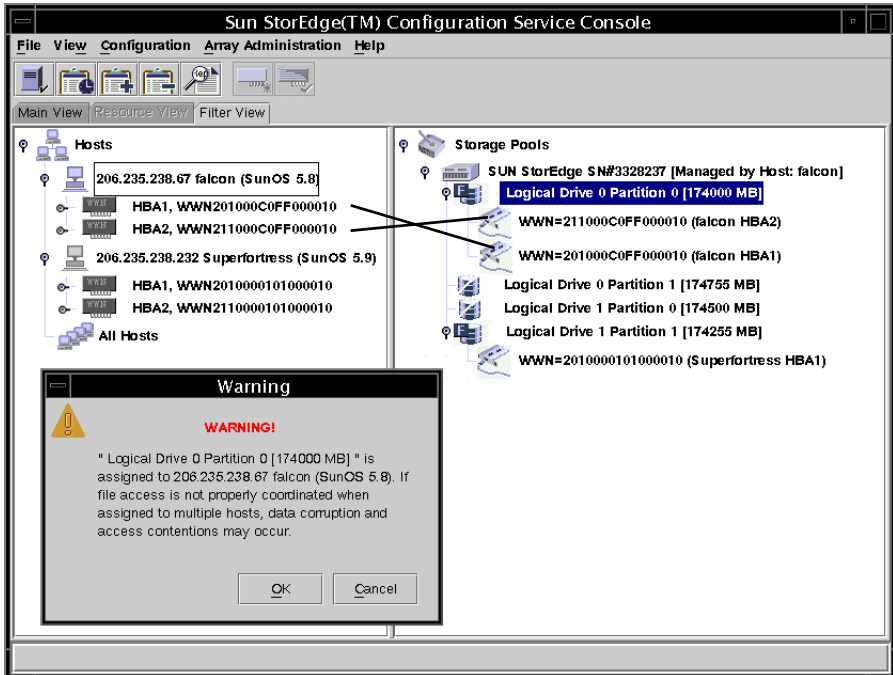
The assignment of the primary controller or secondary controller to an HBA card is selected for the initial LUN filter assignment, and cannot be changed unless you delete the LUN filter assignment and start over.



To apply a LUN filter to a second host for one array, repeat Step 6 under [“To Add HBA Devices Manually”](#) on page 62 through Step 3 under [“To Map Logical Drives to Hosts”](#) on page 65.

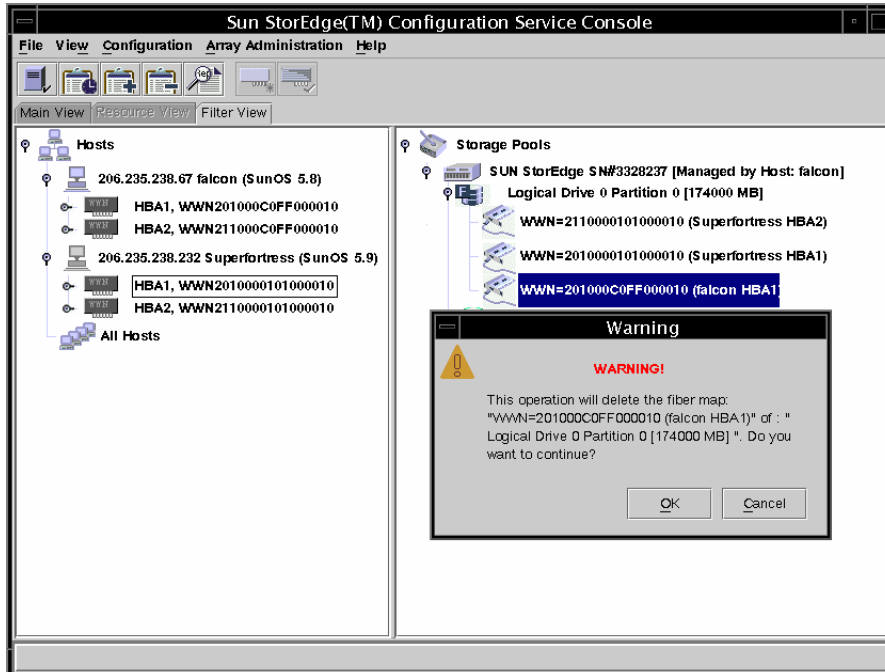
A warning prompt is displayed informing you that two users accessing the same logical drive at the same time may cause data corruption; click OK to add the additional host.

In the following example, the Sun StorEdge 3510 Fibre Channel device has LUN filters (note the letter F for filter) to two separate hosts.



▼ To Delete a LUN Filter

1. Select the associated WWN.
2. Drag and drop it to Storage Pools.
3. When the prompt to confirm the deleted mapping is displayed, click OK.
4. When the second confirmation is displayed, click OK.



Monitoring the Array

This chapter explains how to monitor the array using Sun StorEdge Configuration Service. It describes the main window and the component view windows that are available when an icon is double-clicked in the main window. It also explains how the Event Log works and how to use the reporting function. It is organized into the following sections:





- [“The Main Window” on page 69](#)
- [“Viewing Detailed Device Information” on page 76](#)
- [“Event Log” on page 94](#)
- [“Save Report” on page 98](#)
- [“View Report” on page 101](#)
- [“Out-of-Band Storage Management” on page 101](#)
- [“Managing Storage Through the Web” on page 107](#)



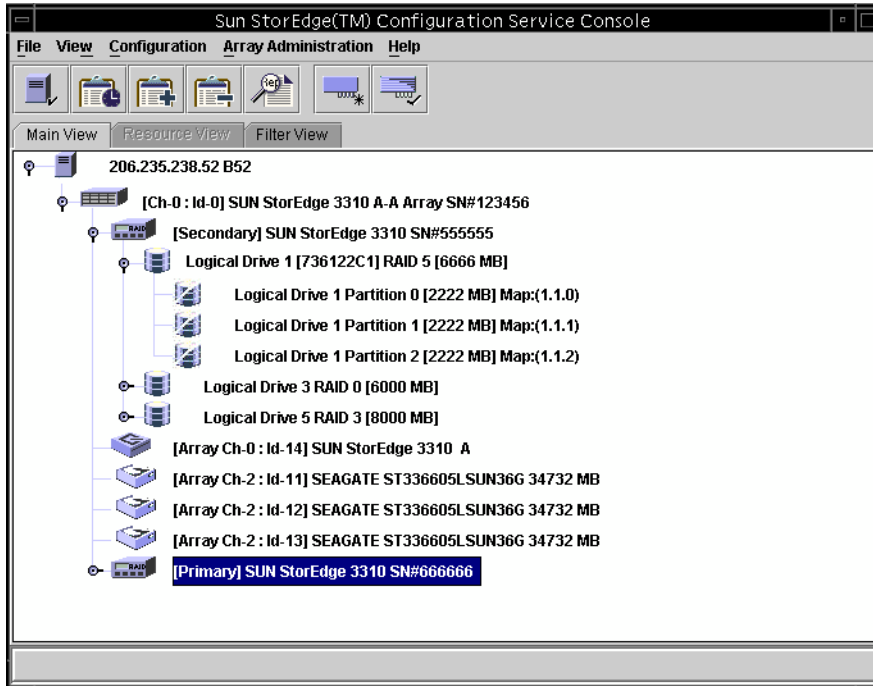
Caution – Sun StorEdge Configuration Service can monitor and manage up to 32 arrays at one time. However, console response time can decrease as the number of arrays increases.

The Main Window

The main window provides *status at a glance* for all array devices. It indicates the status of the devices and logical drives connected to servers monitored by the console. The devices include array controllers, disk storage enclosures, physical disk drives, and other SCSI devices.


This window's tree structure offers detail for all devices connected to each server. The  or  container symbol at the left side of the tree indicates whether the display of devices is expanded or collapsed. The  container symbol means that you can click to display more devices. The  container symbol indicates that all devices at and below that level are shown.

The following figure shows an example of the expanded view of the main window.



Note – The SN# represents the unique ID of the array.




The following figure shows an example of a collapsed view.



For more information, you can double-click an icon to open a component view window. The component views are covered later in this chapter.

Another feature of the main window is that device states are color-coded and symbol-coded so that it is easy to identify when a device is in a state that requires attention. The status is propagated along the device tree, enabling you to trace a failure down to the device level. See [TABLE 6-1](#) for device status details.

TABLE 6-1 Device Status State

Color	Symbol	State
Purple	None	The group, server, or device is online.
White	None	The user has not logged into this server.
Yellow		One or more components of this group or server are not working properly, but the array is still functional. For example, a logical drive that has one failed physical drive is operating in a degraded state. For the Sun StorEdge 3510 FC array, it might also indicate that the battery is going to expire in 21 days (see “Battery Information” on page 89 for more information), or that the in-service date has not been set for a replacement battery. (Refer to the <i>Sun StorEdge 3000 Family FRU Installation Guide for 2U Arrays</i> for details on installing a replacement battery and “To Verify the In-Service Date When Replacing a Battery” on page 91 for information on the in-service date.)
Red		One or more components of this group or server are not working. A logical drive with two failed disk drives or an enclosure with three failed fans would be in a critical state. For the Sun StorEdge 3510 FC array, it might also indicate that the battery has expired. (See “Battery Information” on page 89 for more information.)
Gray		The group, server, or device is unresponsive.

Online Help


To access online help, choose Help → Contents. The online help is in HTML format and depending on the OS, can be run through Microsoft Internet Explorer or Netscape Navigator™. It includes information about major functions within the program.

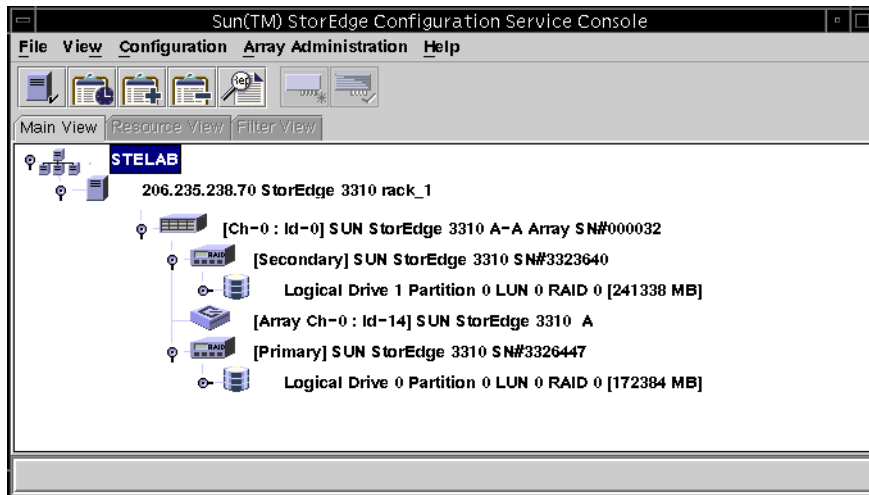
Tree View of Product Configurations

For each server (or group of servers), devices in the tree view are arranged in a hierarchical order with the server (or group) at the top, followed by the array controllers. How the remaining devices — logical drives, physical drives, and enclosures — are shown might vary slightly, depending on which array and controller are being displayed. This section provides window views that might be depicted in the tree view.

Groups

A *group* is a logical collection of servers within the device tree. This new data object enables multiple servers to be contained under a single category.

The group object  looks and behaves like all tree objects. It has an iconic representation and generally assumes the status of its servers. The following figure shows the group icon expanded to show the subordinate or underlying server objects.



A group is a new data object that enables multiple servers to be contained under a single category. Groups are similar in concept to domains. They enable you to organize servers. Instead of a linear tree of all the managed servers, you can organize the servers into like sets or groups.

Groups are color-coded and symbol-coded similarly to servers. States with their corresponding colors have the following order of precedence:

- Critical – Red (the highest)
- Nonresponding – Gray

- Degraded – Yellow
- Optimal – Purple
- Not logged into – White

A server icon assumes the color of the highest state of any of its storage system. Likewise, a group icon takes on the highest state of any of its servers, with the following exceptions involving non responding or unlogged servers.

When a group icon is disabled (inactive), it indicates that all the servers attached to that group are not responding; if any number less than all of the servers in the group is not responding, the group icon appears red, representing a critical state. For example, if there are four servers in the group and three or fewer servers are not responding, the group icon is color-coded red.

When a group icon is color-coded white (not logged in), it indicates that one or more servers within that group have not yet been fully configured, or could represent a period of status transition. TABLE 6-2 gives an example of the different color coding of a two-server group.

TABLE 6-2 Two-Server Group Colors

Server 1 Icon Color	Server 2 Icon Color	Group Icon Color
Gray	Gray	Gray (not responding)
Gray	Yellow, Red, Purple, or White	Red (critical)
White	White	White (not logged into)
White	Yellow	Yellow (degraded)
White	Red	Red (critical)
White	Purple	Purple (optimal)
Red	Any Color	Red (critical)
Yellow	Yellow or Purple	Yellow (degraded)
Purple	Purple	Purple (optimal)

Groups are not required. You can configure the program for no groups and fifteen servers, for example, or for one group with ten servers underneath, with an additional five at the top level. The program enables any combination.

The number of groups permitted and the number of servers within a group is limited only by available system memory. If a server is a member of a group and a user deletes that group from the group list box, servers are reassigned in that group to the no group category. The tree is remapped in the main window.

The Monitoring Process

The console monitors storage devices on a network by communicating with an agent on the servers.

When the program is started, the console software begins by establishing contact with the agent on each managed server if Auto Discovery was specified when the server was configured. If Auto Discovery was not specified, you have to double-click each server and provide a password to start the discovery process for that server.

Establishing a TCP/IP connection between the console and the agent on each server and receiving back inventory information can take several minutes, depending on the complexity of the network. When this is occurring, the server icon in the main window displays a satellite dish icon on its right side. Once the inventory is refreshed, the satellite dish symbol is replaced by an active server symbol.

The agent on each server does a periodic scan of its inventory to check for changes. If there is a change, the agent sends an event to the console. Depending on the event, the console might request the inventory from that server's last scan to use for updating the representation of the server as displayed on the main window. During this process, the satellite dish icon is attached to the server icon, and you cannot do any configuration and array activity commands on that server until the refresh process is completed and the console main window is updated.

When the program is running and the server agent and console are connected, the agent pings or transmits a periodic signal to the console to verify the server's status. If the console does not receive a certain number of consecutive responses (handshakes) from the agent, the console marks the server as offline and disconnects from the server. The server's icon is made inactive and adjacently marked with a question mark symbol.

If a nonactive server was originally made active through the Auto Discovery function, the console periodically tries to reestablish server communications.

Occasionally, you might want to have an inventory performed on a server between the periodic status scans. To do this, double-click the server's icon to display the Server View window, and then click Rescan in that window.

Auto Discovery Options

If you have selected the Auto Discovery option (during the process of adding servers to the Managed Servers list, see [“Select or Deselect Automatic Discovery of servers:” on page 13](#)), the program automatically scans and performs an inventory on these servers. You are not required to provide a monitoring password to retrieve information provided by the program. Depending on the complexity of the network and the number of servers, it can take several minutes for the discovery process to finish.

However, if you choose not to use the Auto Discovery option on startup, the servers’ icons are white, indicating that there is currently no information available for these servers. In this case, you need to double-click each server icon and specify the appropriate monitoring password.

You can also choose File → Login. Once the password is specified for a server, the discovery process begins to perform an inventory on the selected server.

Note – If you need a very secure environment where access to server inventory data is restricted, select No for Auto Discovery (see [“Select or Deselect Automatic Discovery of servers:” on page 13](#)).

Viewing Detailed Device Information


Sun StorEdge Configuration Service provides detailed information about each agent or server and about the array devices that are connected to it. The fields on each of these views vary, depending on the capabilities of the disk controller.

Except for Save Report and View Report, which are located under the File menu, the commands and windows described in this section are accessed through the View menu.

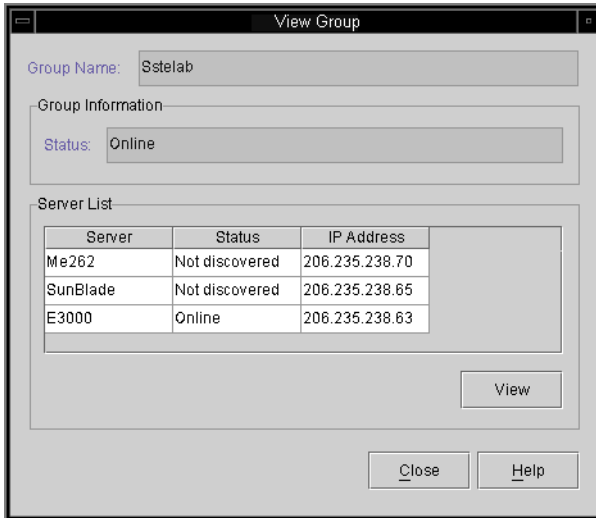
- [View Group](#)
- [View Server](#)
- [View Controller Configuration](#)
- [View Logical Drive](#)
- [View Physical Drive](#)
- [View Enclosure](#)
- [Double-click the enclosure icon.](#)
- [Array Administration Progress](#)
- [Agent Options Management](#)
- [Save Report](#)
- [View Report](#)

View Group

View Group displays the servers that make up the group that is selected in the main window.

To access View Group, double-click the group icon  in the main window or select the group icon and choose View → View Group.


The Server List displays all servers that are attached to the specified group.

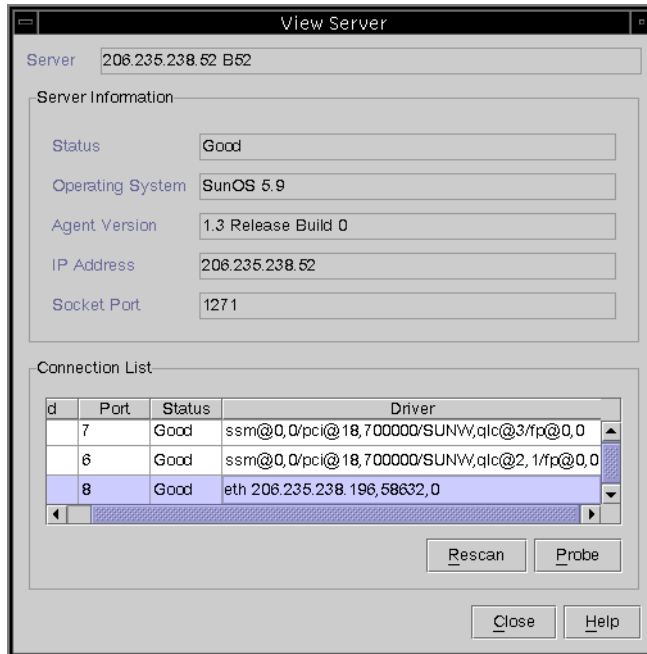


To view more information about any of the servers listed, select the appropriate status record in the list and click View, or double-click the status record. A window describing the corresponding server is displayed.

View Server

View Server displays the characteristics of the server that is selected in the main window.

To access View Server, double-click the server icon  in the main window or select the server icon and choose View → View Server.



Note – For SCSI and FC arrays, dual-port HBAs are displayed as separate entries because they have separate channels to which devices can be connected.


Part of the TCP/IP network, the socket port makes a connection between the server and the client.

- Connection List – Displays host adapters and array controllers that are installed in or connected to the specified server.
- Rescan – Sends a command to the selected server to rescan existing inventory and refresh the console. Normally, each managed server performs a periodic scan of its inventory and updates the console if there is a change. If you want to perform a diagnostic of all assigned devices, click Rescan; selecting this button essentially circumvents the periodic and automatic status updates.


- Probe – Sends a command to the selected server to probe for new inventory (for example, RAID controllers, JBODs, and expansion units). Whenever a new device has been added, or a device name has changed and you want it to be displayed in the main window tree view immediately, click Probe.

View HBA Card

View HBA Card displays the characteristics of the Host Bus Adapter (HBA) card that is selected in the main Configuration Service window. It is only displayed when using out-of-band management.


To access View HBA Card, double-click the host adapter icon  in the main Configuration Service window or select the host adapter icon and select View → View HBA Card.

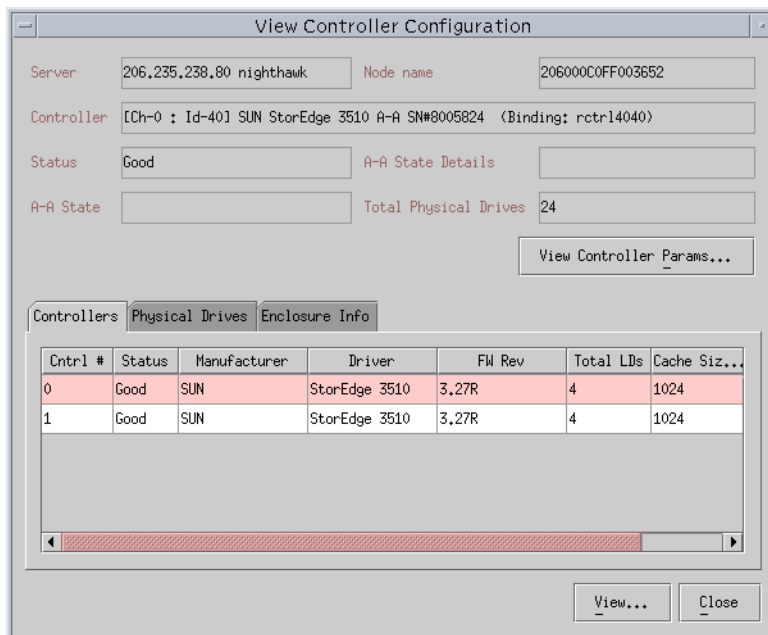
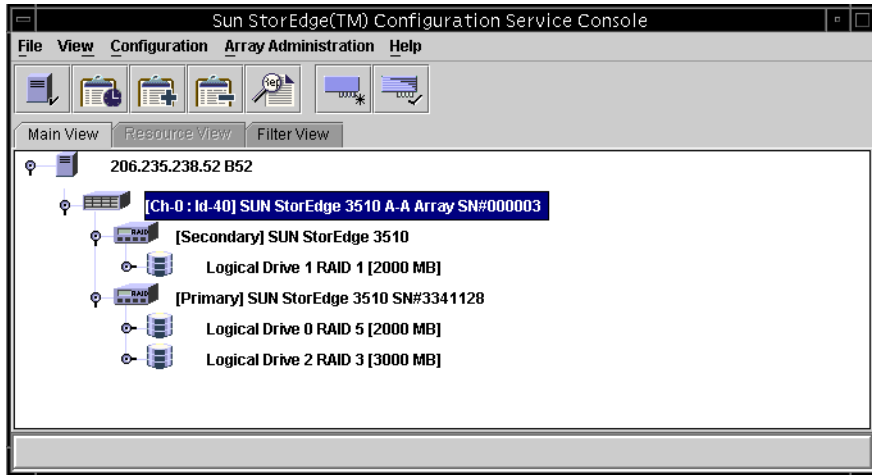
View HBA Card displays the status of the host adapter card and its device driver and provides a list of the devices connected to the adapter. Configuration Service channels start with number 0. The server listed at the top of the window is the server in which the host adapter is installed.



View Controller Configuration


View Controller Configuration displays the components that make up the array.

To access View Controller Configuration, double-click the array icon  in the main window, or select it and choose View → View Controller.




The tabs at the bottom section of the window provide detailed information about the RAID controller's LUNs, on the physical drives attached to it, on the enclosure where the controller is located, and on the configuration of the peripherals. To bring the information for another tab into view, click the tab.

Controllers Tab

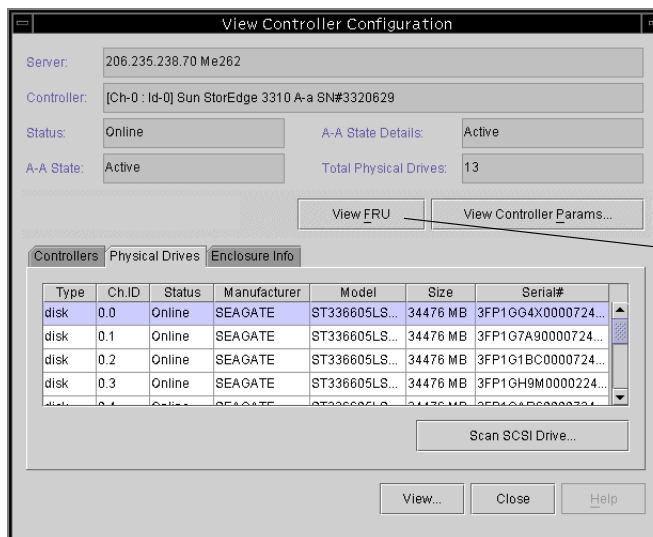
To access this window, double-click the array icon  in the main window, or select it and choose View → View Controller. Then click the Controller tab.

The Controller tab lists the controllers. To view more information about the controllers, double-click the desired controller or select the controller and click View. The SCSI and FC arrays have primary and secondary controllers. Depending on whether the controller is primary or secondary, the View Primary or View Secondary Controller Configuration window is displayed. See [“View Primary/Secondary Controller Configuration” on page 83](#).

Physical Drives Tab

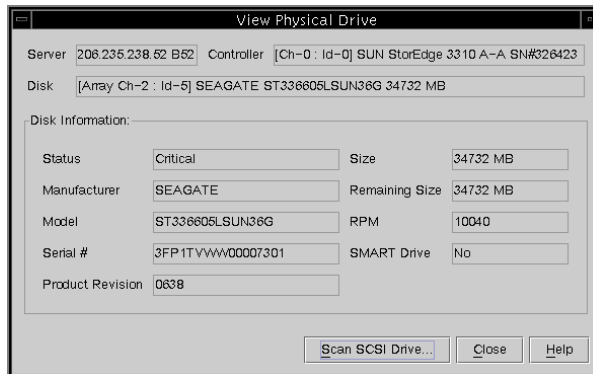
To access this window, double-click the array icon  in the main window, or select it and choose View → View Controller. Then click the Physical Drives tab.

The following figure shows the View Controller Configuration with the Physical Drives tab displayed.





View FRU is not displayed for the Sun StorEdge 3510 Fibre Channel array.

The Physical Drives tab lists the physical drives associated with the array. To see details about any of the physical drives listed, double-click the drive or select the drive and click View. The View Physical Drive window is displayed. If a drive fails, the Scan SCSI Drive button is displayed.



For more information about View Physical Drive, see [“View Physical Drive” on page 87](#).

Enclosure Info Tab

To access this window, either double-click the Enclosure icon  in the main window, or select the array icon  and choose View → View Controller. Select the Enclosure Info tab. Select an enclosure and click View.

The information in the Enclosure Info tab includes the status of the power supplies, fans, battery, and the enclosure temperature. For more information about View Enclosure, see [“View Enclosure” on page 88](#).

View FRU

For the Sun StorEdge 3310 SCSI array only, to display the field-replaceable unit identification (FRU ID) information for the array, click View FRU. For example FRU ID information, see [“Double-click the enclosure icon.” on page 92](#).

View Controller Parameters

Click View Controller Parameters to view detailed information about Channels, RS 232, Cache, Disk Array, Drive I/F, Host I/F, and Redundancy. To view the information, click the tab that corresponds to the item you want to view. To change controller parameters, see [“To Change Controller Parameters” on page 153](#).

Server: 206.235.238.53 ultra5 Controller: 1-Id-3j SUN StorEdge 3310 A-A SN#900000

Controller Data

Firmware Version: 3.25N Boot Record Version: 1.31G

Serial Number: 900000 CPU Type: PPC750

Controller Name: Cache Size: 512MB ECC SDRAM

Controller Unique ID: 112233 Append File ID: 3310 325N v1.19

Channel: RS 232 Cache Disk Array Drive I/F Host I/F Redundancy Network

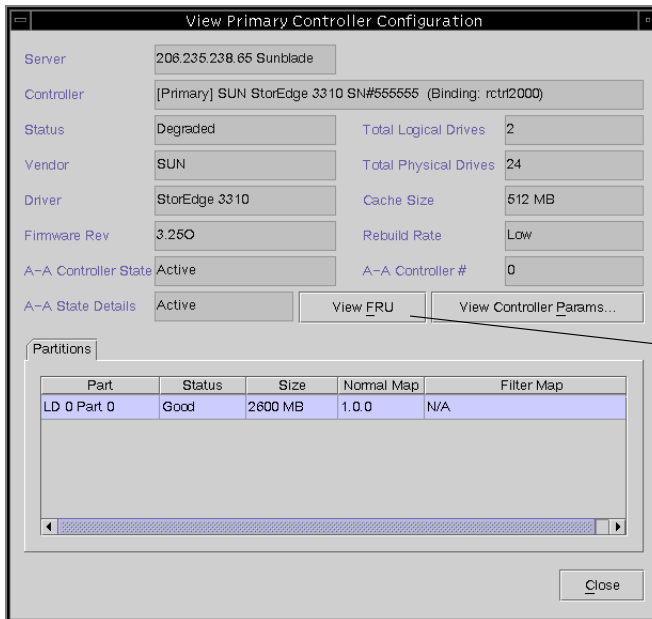
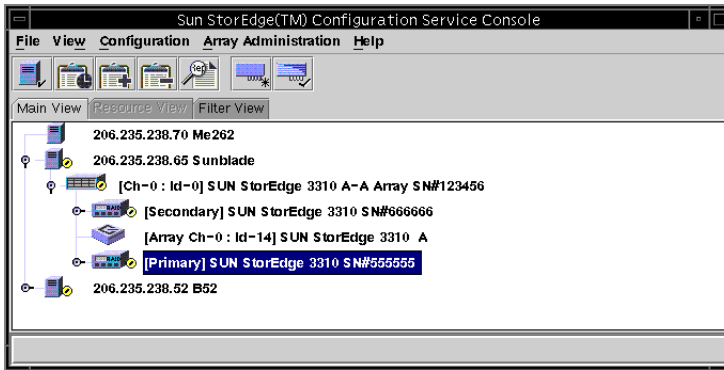
Chl#	LogChl.	Mode	Type	PID	SID	DefClk	DefWid	Term	CurClk	CurWid
0	0	Drive	SCSI	6	7	80 MHz	Wide	Off	80 MHz	Wide
1	0	Host	SCSI	0	1	80 MHz	Wide	Off	80 MHz	Wide
2	1	Drive	SCSI	6	7	80 MHz	Wide	Off	Async	Narrow
3	1	Host	SCSI	3	2	80 MHz	Wide	Off	80 MHz	Wide
6	0	N/A	RCCOM	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Close

View Primary/Secondary Controller Configuration

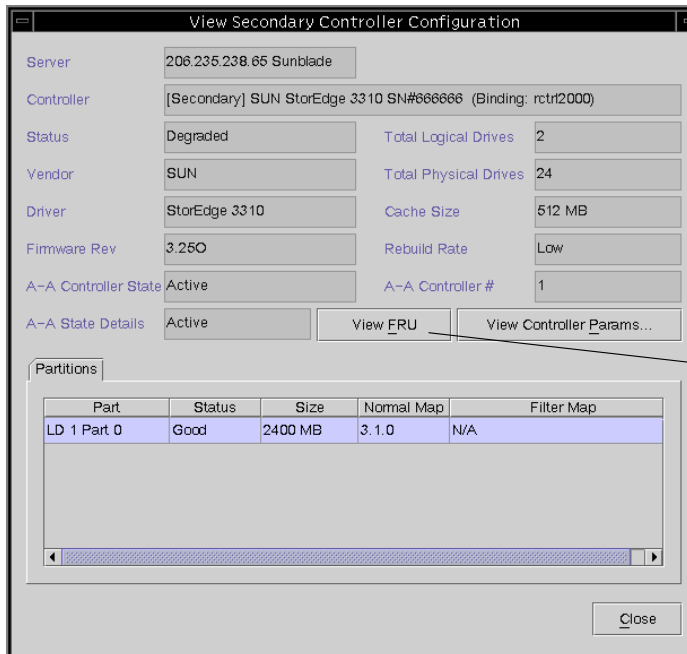
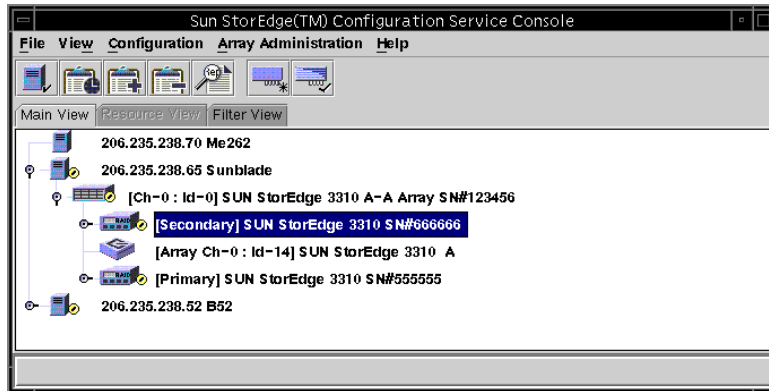
View Primary and View Secondary Controller Configuration windows display the characteristics of the controller that is selected in the main window. To display the primary controller characteristics, select the primary controller from the main window and choose View → View Controller, or double-click the primary controller.

The following two examples show the primary controller.





View FRU is not displayed for the Sun StorEdge 3510 Fibre Channel array.

The following two examples show the secondary controller.



View FRU is not displayed for the Sun StorEdge 3510 Fibre Channel array.

Logical Drives Tab

To access this view window, double-click the primary or secondary controller icon  in the main window, or select the controller icon  and choose View → Controller.



Click View Controller Params for editable component parameters on the RAID Controller Parameters window.

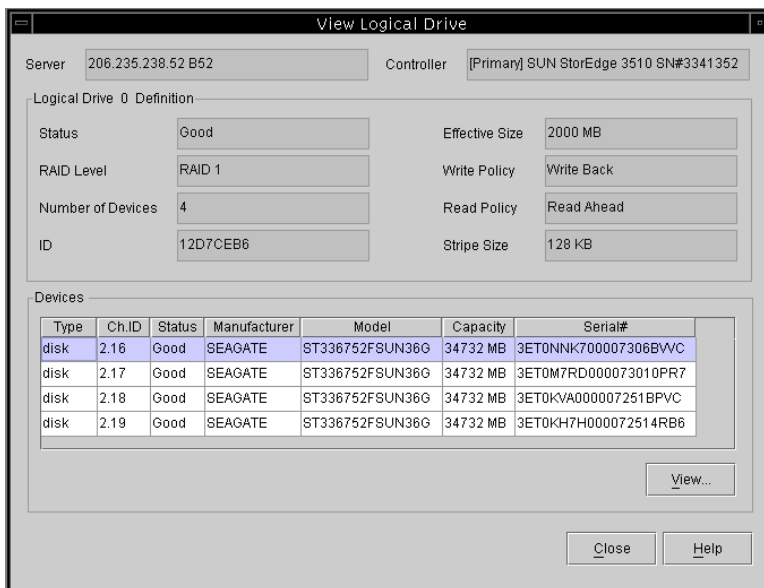
To view a logical drive, select and double-click a drive from the Logical Drives tab, or select a drive and click View.

View Logical Drive

View Logical Drive displays the characteristics of the logical drive that is selected in the main window.

To access this view window, use one of the following methods.




- Double-click the logical drive icon  in the main window.
- Select the logical drive icon and choose View → View Logical Drive.
- Double-click the controller icon , select a logical drive in the list box, and click View.



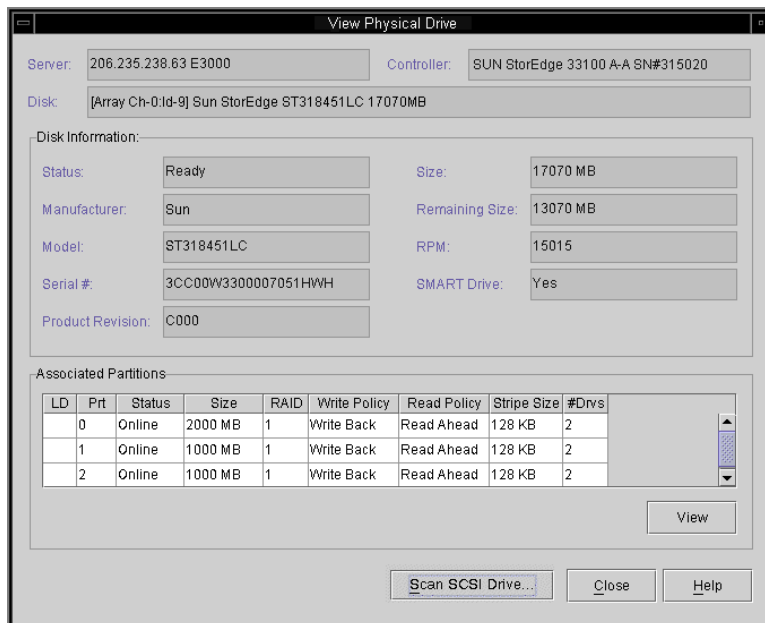
For more information about any of the disk drives listed, either double-click the drive's status record, or select the drive and click View. The View Physical Drive window is displayed.

View Physical Drive

View → View Physical Drive presents the characteristics of the selected physical device. The title of the view window that appears varies, depending on the device selected.

- For a disk drive , it is called View Physical Drive window.
- For a tape drive , it is called View Tape window.
- For a CD-ROM , it is called View CD-ROM window.

You can access one of these windows by double-clicking the physical device in the main window, or by selecting the device and choosing View → View Physical Drive.



- **Associated Partitions** – Displayed only when the selected device is associated with one or more logical drive partitions.
- **View** – Active when there are associated logical drives
- **Remaining Size** – Remaining unused capacity when part of a disk drive's capacity has been used in one or more logical drives
- **RPM** – The revolutions per minute of the physical drive

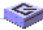
- SMART Drive – Indicates whether the drive has predictive failure capability

You can also select the physical drive and select View. The View Logical Drive window is displayed without the table.

The View Tape and View CD-ROM windows are similar to that shown in the previous figure, but the categories differ slightly.

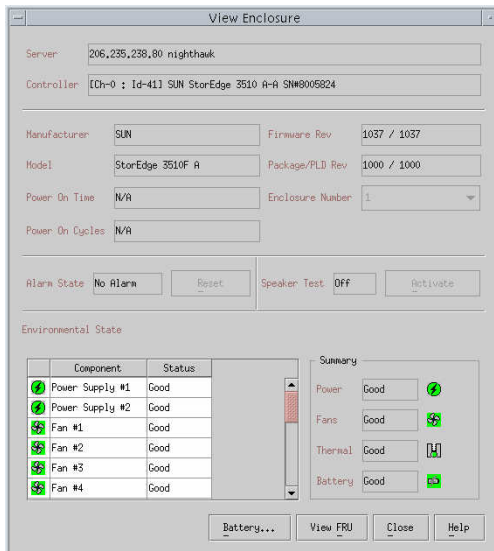
View Enclosure

The View Enclosure window displays the component and alarm characteristics of an enclosure that is selected in the main window. (For SCSI, the EMU is identified by the name Sun StorEdge 3310 A; Id is always 14 or 15. For Fibre Channel, the SES is identified by the name Sun StorEdge 3510 A; Id is always the last Id within the enclosure in which the SES is contained.)

1. Select an EMU (SCSI array) or SES (Fibre Channel array) icon .
2. Choose View → View Enclosure.

The upper two sections of the window identify the enclosure and provide related information. Note that when there are multiple enclosures, you can use the Enclosure Number list box (located in the upper right corner of the window) to reflect another enclosure attached to the same controller.

For the 3510 FC array, to display battery information, click Battery and see [“Battery Information” on page 89](#). To display the FRU IDs and information for all FRUs in the array, click View FRU and see [“View FRU” on page 92](#).



The Alarm State section of the window applies only if you have an array with a SCSI Accessed Fault-Tolerant Enclosure (SAF-TE) (for a SCSI array) or SCSI Enclosure Services (SES) (for a Fibre Channel array) card.

- If the Alarm State field displays `Alarm!`, it means that there is an alarm state caused by a failure of a drive, fan, power supply, or battery, or by an abnormal temperature in the enclosure. The audible alarm on the box sounds.
- For the program to continue monitoring after this alarm condition starts, you must push the Reset button on the right ear of the array.

Note – The Reset button on the View Enclosure window is reserved for future use.

- Only users assigned configuration security privileges are allowed to Reset the alarm through the program. Replace the failed component or check the enclosure if its internal temperature is abnormal.

Note – Controller events can also cause an audible alarm. Pushing the Reset button has no effect on audible alarms caused by a controller event. See [“To Mute the Controller Beeper” on page 167](#) for information about muting the beeper.

The lower section of the window provides status of the components in the window. When there is a device in the slot, it is represented by a drive icon with a color and a symbol representing its device state.

Battery Information

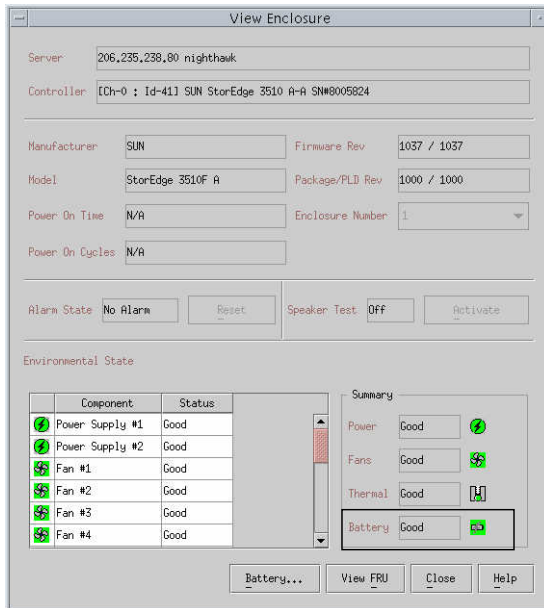
Note – The Battery Information window does not apply to the Sun StorEdge 3120 SCSI array or the Sun StorEdge 3310 SCSI array.

In the event of a power failure, the battery maintains power to the cache for 72 hours. When power is restored, the data in cache is dumped to disk. For the Sun StorEdge 3510 FC array, Sun StorEdge Configuration Service monitors the usable life of the battery and displays its status in the Battery Information window. The program calculates the battery expiration date using the battery type, manufacture date, and in-service date, which have been programmed at the factory.

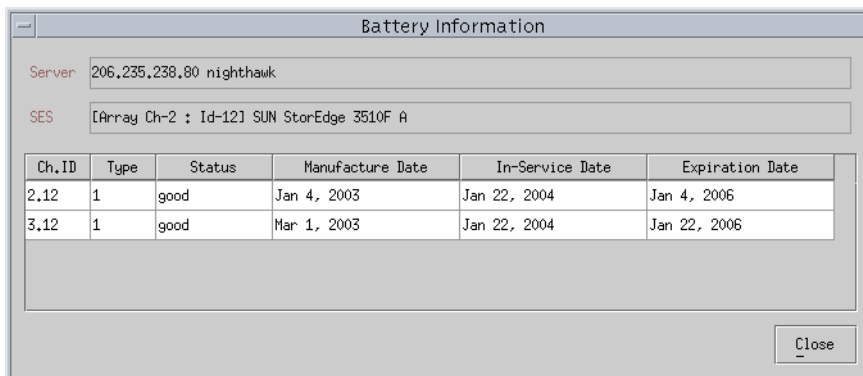
Note – For a battery FRU, you need to verify the in-service date so that Sun StorEdge Configuration Service can set it as explained in [“To Verify the In-Service Date When Replacing a Battery” on page 91](#).

The enclosure icon on the main window displays a degraded (yellow) status 21 days before the battery is going to expire. The enclosure icon also displays a warning (yellow) status if the in-service date has not been set for a battery FRU. A critical (red) status is displayed when a battery has expired. See “[Device Status State](#)” on [page 71](#) for device status symbols.

To view the battery status, choose View → View Enclosure or double-click the enclosure. The View Enclosure window is displayed, showing the battery status in the Summary box.

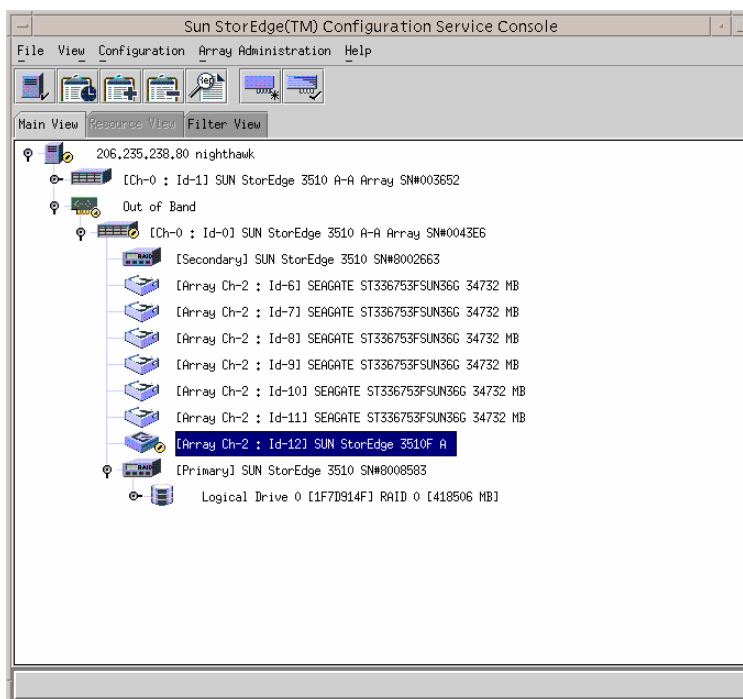


To view battery information, including type, status, manufacture date, in-service date, and expiration date, click Battery. The Battery Information window is displayed.




▼ To Verify the In-Service Date When Replacing a Battery

When Sun StorEdge Configuration Service detects a battery FRU, the enclosure icon displays a degraded (yellow) status symbol as shown in the following example.



1. Double-click the enclosure icon.

Sun StorEdge Configuration Service calculates the battery expiration date using the in-service date (date that the battery is put into service), which is based on the host clock. The program prompts you to verify the date by displaying the following message:



2. If the host clock is correct, click Yes.

The following confirmation message is displayed. Click OK.



Sun StorEdge Configuration Service sets the in-service date and displays the date in the In-Service Date field in the Battery Information window.

3. If the host clock is incorrect, click No and reset the clock so that Sun StorEdge Configuration Service can prompt you to verify it again and set the in-service date.

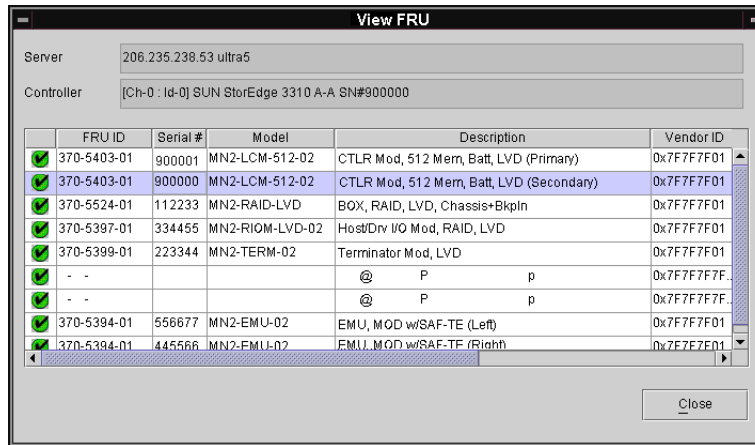


Caution – If you do not reset and verify the in-service date, Sun StorEdge Configuration Service cannot accurately calculate the battery expiration date.

View FRU

A FRU is a field-replaceable unit. It is a part used to assemble a new system or to repair a system in the field. The Sun FRU ID (field-replaceable unit identification) program is a Sun solution for capturing, transmitting, and analyzing FRU-specific configuration, diagnosis, and failure information residing on the FRU.

Choose View → View FRU to display the FRU IDs and information for all FRUs in the array including Serial Number, Model, Description, Vendor ID, Time (time the FRU was programmed), and Location.



Note – You can also view the FRUs for a specific controller for the Sun StorEdge 3310 SCSI array by clicking View FRU on the View Controller Configuration window, the View Primary Controller Configuration window and the View Secondary Controller Configuration window.

Array Administration Progress

Array Admin Progress displays the progress of new logical drive(s) initialization. This command is accessed by choosing View → Array Admin Progress.

Agent Options Management

Agent Options enables you to customize the agent options, including polling time, periodic device discovery time, and SMART monitoring.

To access Agent Options, choose View → Agent Options Management. For more information see, [“To Configure Agent Parameters” on page 22.](#)

Event Log

The console receives, logs, and displays events generated by managed servers and by the console itself.

The majority of the events are generated by the agents on the managed servers and occur when there are:

- Status changes on any device on a managed server, including the server itself (because of device failure, malfunction, or disconnection)
- Inventory changes (addition or removal of devices)
- Configuration changes (initial configuration setup and subsequent configuration changes)
- Array processes running on the server (initialization, parity checking, rebuilding)

Although array processes are initiated by the console, it is the server agent that generates operation notification events after these processes start on the server. For details about full event monitoring and email notification capabilities including sending email messages to a specified list when errors occur, see [“Email and SNMP” on page 207](#).

The console generates a much smaller number of events. For example, it generates an event if it does not receive a certain number of consecutive heartbeats from a managed server.

When the console receives any event, it logs it in to the Event Log file, `eventlog.txt`, and displays it in the Event Log window. Also, if the event occurs on a server, the notification of the event is sent to that server’s OS event log. On an NT server, it would go to the NT event log. In addition, when the event occurs on a server and that server is set up to send traps to an SNMP enterprise management console, such as HP OpenView, the server agent also sends a trap message to that computer.

Depending on the event received, the console might initiate a refresh process to request the inventory from the last periodic scan of the server involved, so the console can update the server’s inventory on the main window.

During this refresh process, the satellite dish icon is attached to the server icon, and you cannot perform any configuration and array activity commands on that server until the process is completed and the main window is updated.

Event Log File

The Event Log window displays up to 500 events at a time. If you have more than 500 events, only the most recent 500 are displayed in the Event Log window; however, the program does not delete any events from the Event Log file, `eventlog.txt`, until more than 10,000 events have been logged.

- After 10,000 events, the program reduces the Event Log file to the most current 500 events and then accumulates events until the limit of 10,000 is exceeded again.
- The fields of each event record are separated by a semi-colon so you can easily import the file into a database.
- `eventlog.txt` is located in the directory where the console program files are installed.

Note – If the event log appears not to contain all of the events from the managed array, close and reopen the console.

The events from the agent are logged into the system log of the host where the agent is installed, even if the console is not running. The following table lists the locations where the events are logged to in each OS.

TABLE 6-3 Event Log Location

OS	Event Log Location
Solaris OS	<code>/var/adm/messages</code> (Also shown on the console)
Linux OS	<code>/var/log/messages</code>
Microsoft Windows OS	The application log of the system, which can be viewed using Event Viewer. You can also read the event log directly from the file <code>\Program Files\Sun\sscs\eventlog.txt</code>
HP-UX OS	<code>/var/adm/syslog/syslog.log</code>

▼ To Write Events to a Log File for an IBM AIX Host

For an IBM AIX OS, the event logs are not logged by default. You might need to change `/etc/syslog.conf` to enable it to write to a log file.

1. **Modify `/etc/syslog.conf` to add the following line:**

```
*.info /tmp/syslog rotate size 1000k
```

2. **Make sure the file that is specified in the added line exists.**

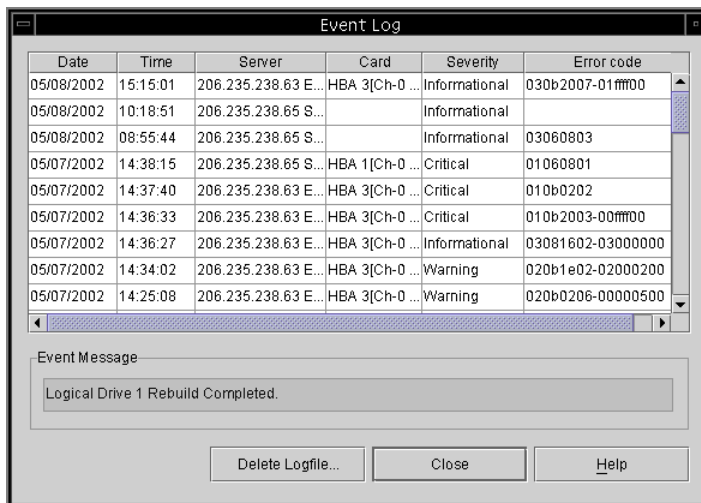
If it does not exist, you need to create it. For example, in the above configuration, you would create a file named `/tmp/syslog`.

3. **Change to `/tmp/syslog` and restart the `syslog` by typing:**

```
kill -HUP `cat /etc/syslog.pid`
```

Event Log Window

To access the Event Log window, choose `View → Event Log`. You can hide this window by clicking `Close`. You can then reopen it (from the `View` menu) without losing any content.



The consoles begin to receive events when they are running, regardless of whether the Event Log window is open.

1. To delete the log file, click Delete Logfile.

The Confirmation window is displayed, prompting you to save the log file.

2. Select one of the following options:

- Select yes at the prompt, select a folder and a file name, and save the log file.
- Select no at the prompt.

The contents of the log file is deleted.

Note – You can also save and delete the contents of the `eventlog.txt` file using the Save Event Log and Delete Event Log icons on the toolbar.

Each event record contains the fields shown in the following table.

TABLE 6-4 Event Record Fields

Date	The date on the server when the event occurred.
Time	The time on the server when the event occurred.
Server	The IP address of the server and the server name.
Card	The card name, if applicable, for the event.
Severity	One of three severity levels: Critical, Warning, or Informational. (These states are described in the following subsection.)
Error Code	The basic error code and the extended error code, separated by a dash.
Text Message	A text message describing the event.

Severity Levels

- Critical – A message that does require intervention by the network administrator, such as the failure of a device, power supply, or fan.
- Warning – Warning messages generally indicate internal program events. However, if you see a large number of these messages, it might mean that there is a problem with the server or the network.
- Informational – A message about the devices on the server that does not require intervention by the network administrator.

You receive alarm forwarding for the level selected and any other levels of a higher severity. Thus, if you choose Informational, you are also notified of all alarm conditions. However, if you choose Critical, only Critical alarms are received.


For further information about messages, see [“Troubleshooting” on page 223](#).

Save Report

The Save Report option creates a text file containing all the information available to the program about a specific array.


1. Choose File → Save Report.

The Report dialog box for confirming the server’s ID is displayed. The Exportable to spreadsheets option enables you to save the report with delimiters (Comma, Semicolon, Tab, Colon, and Vertical Bar) for export to spreadsheets.



2. Select Save report.

The Save Report File window is displayed.



3. Type a file name to identify the file and click Save.

The program writes the report on the inventory and status of the selected server and its devices.

The default save location for the report file is in the installation directory and the default file extension is `.xml`. It might be helpful to create a subdirectory for saving reports, so they do not clutter up the installation directory.

A report includes the following information:

- General array details
 - Array name, status, manufacturer, model
 - Firmware version
 - Boot record version
 - MAC, IP, netmask, and gateway address
 - Serial numbers of controllers
- General overview of the array configuration
 - Total number of logical drives, logical volumes, and partitions
- Controller parameters
 - Host and drive channel settings and IDs
- Detailed list of each logical drive
 - RAID level
 - Number and size of physical drives in a logical drive
 - Number and size of partitions per logical drive and their mappings
- SAF-TE/SES information
- List of FRUs
- Details of the physical (hard) drives

The following excerpt is from an example report for a Sun StorEdge 3510 Fibre Channel array in .xml format, which can be used as input into another program.

CODE EXAMPLE 6-1 Excerpt From Sun StorEdge 3510 Fibre Channel Array Report


```
<raidbaseview>
  <raidsystem>
    <name> SUN StorEdge 3510 A-A Array SN#000198 </name>
    <status>Good</status>
    <manufacturer>SUN</manufacturer>
    <model>StorEdge 3510</model>
    <firmware_version>3.27K</firmware_version>
    <bootrecord_version>1.31H</bootrecord_version>
    <mac_address>00:C0:FF:00:01:98</mac_address>
    <ip>206.235.238.198</ip>
    <netmask>255.255.255.0</netmask>
    <gateway>206.235.238.1</gateway>
    <primary_sn>3341275</primary_sn>
    <secondary_sn>3341258</secondary_sn>
    <controller_name>198</controller_name>
    <unique_id>198</unique_id>
    <id_of_nvram_defaults>327K 3510 v2.39</id_of_nvram_defaults>
    <total_logical_drives>8</total_logical_drives>
    <total_logical_volumes>0</total_logical_volumes>
    <total_partitions>278</total_partitions>
    <total_physical_drives>24</total_physical_drives>
    <total_safte_ses_devices>1</total_safte_ses_devices>
    <cache_size>1024MB ECC SDRAM</cache_size>
    <cpu>PPC750</cpu>
    <battery>Good</battery>
    <node_name>206000C0FF000198</node_name>
    <fru>
      <idx>0</idx>
      <item></item>
      <serial_number></serial_number>
      <revision></revision>
      <vendor_jedec_id></vendor_jedec_id>
      <manufacturing_date></manufacturing_date>
      <manufacturing_location></manufacturing_location>
```

View Report

Use the View Report option to review a report that was created.

1. Choose File → View Report.

The Open dialog box for selecting the report is displayed.



2. Select the report you want to review and click Open.

Out-of-Band Storage Management

The out-of-band storage management capability enables you to monitor and manage arrays over the network using TCP/IP. Unlike in-band storage management (the standard method of storage management for storage), which requires the agent to be running on the server that is physically attached to the storage, out-of-band storage management does not require the agent to be running on the server that is physically attached to the storage. With out-of-band storage management, if the server that is attached to the storage is shut down, monitoring and maintenance is unaffected.



Caution – If you assign an IP address to an array in order to manage it out-of-band, for security reasons make sure that the IP address is on a private network, rather than a publicly routable network. When your array is on a public network, it is susceptible to viruses, worms, and other malware attacks.

The following figures show examples of in-band and out-of-band storage management configurations.

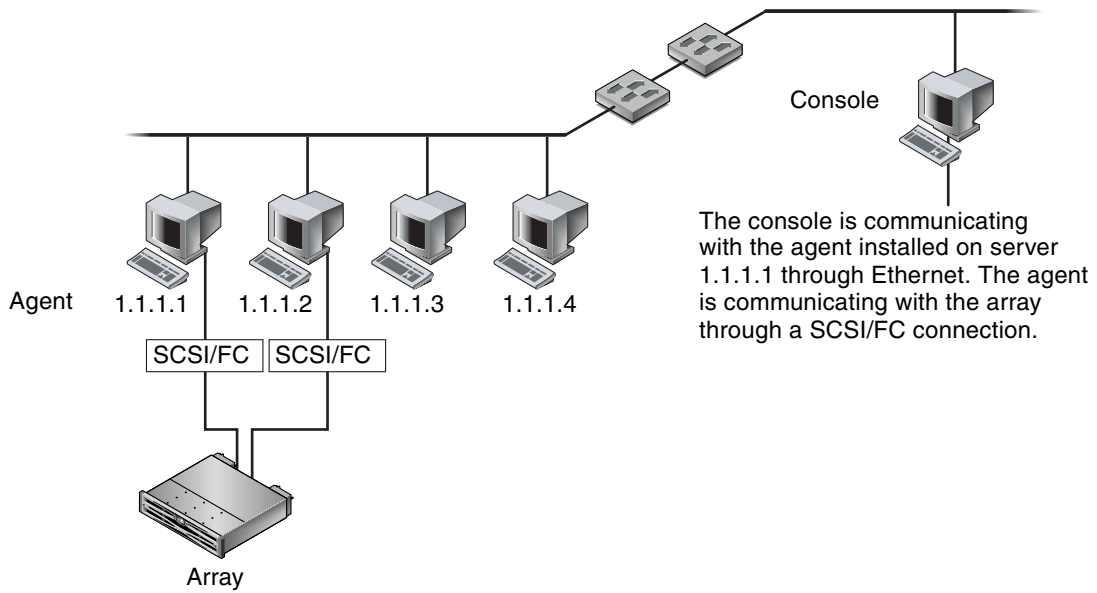


FIGURE 6-1 In-Band Management

For more information about configuring your array for out-of-band management, see [“Email and SNMP” on page 207](#).

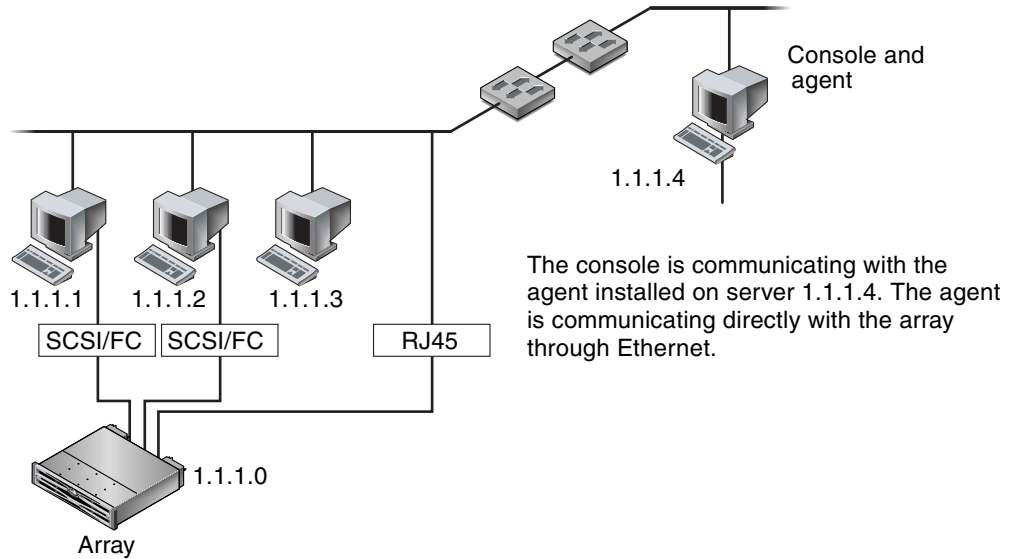


FIGURE 6-2 Out-of-Band Management

▼ To Use Out-of-Band Management

Note – Controller, SAF-TE, and Drive firmware cannot be upgraded through out-of-band management.



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

1. Make sure you have set up a static or dynamic IP address for the array.

If the program has been configured already to manage the array through in-band, you can set the IP address through Change Controller Parameters. To set the IP address, see [“Network Tab” on page 166](#). If the program has not been configured yet, you can set the IP address through an RS-232 terminal. Refer to the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.

2. After setting the IP address, reset the controller.

Choose Array Administration → Controller Maintenance, and then click Issue Reset to the Controller.

3. Select the server.
4. Select View → Agent Options Management.
5. In the Agent Options Management window, check Out-Of-Band Agent preferred over In-Band.

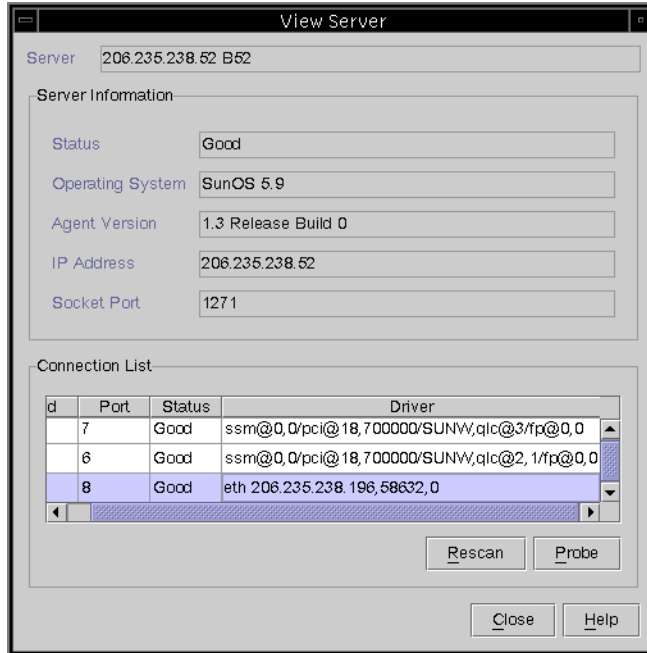
Note – With this option checked, out-of-band becomes the preferred method of storage management. If the out-of-band configuration is removed, the program reverts to in-band storage and monitoring after you start/restart the services.

6. Type the IP address of the array in the IP Address field, and click Add.
7. If you have created a password for the array using the firmware application, type it in the Password field, and then re-type it in the Verify Password field.

Note – By default there is no password set for the array. For information on creating or changing the password, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array.

8. For the program to recognize the out-of-band array and display it in the main window, select the server.
9. Choose View → View Server.

10. Click Probe.

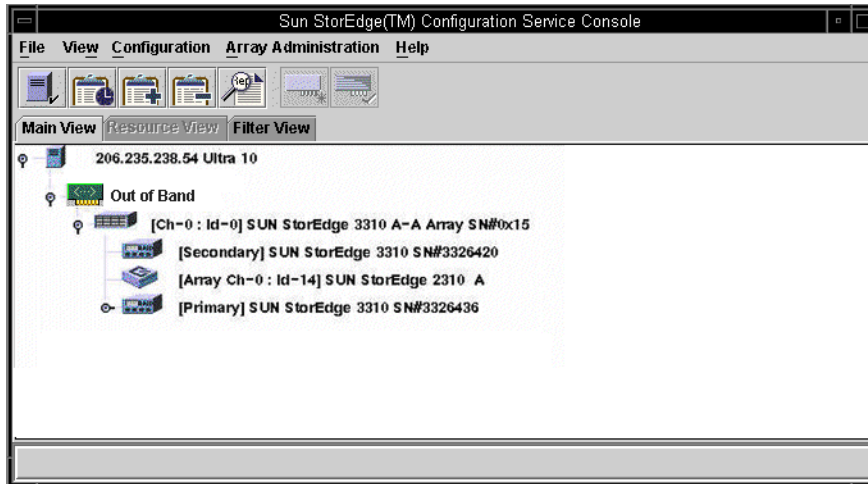


11. If the program has not been configured to manage the array, you need to assign the server to manage the controller.

Choose Array Administration → Controller Assignment. Select a server from the Server to manage this controller list and click Apply.

An out-of-band HBA is displayed in the main window and View HBA Card is displayed in the Menu Bar under the View menu.

Note – If an out-of-band HBA is not displayed in the window, reset the controller.



▼ To Remove an Array From Out-of-Band Management

1. Select the server.
2. Choose View → Agent Options Management.
3. Select the array you want to remove from the Managed Primary Agent list, and click Remove.

The HBA remains displayed in the main window; to remove it, you need to stop/restart services.

Managing Storage Through the Web

The web-based storage management capability enables you to conveniently manage the array through the web without having to load the entire Sun StorEdge Configuration Service package. The following steps describe how to set up the array for web-based management and how to access it through a web browser.

Note – This feature requires Sun StorEdge Enterprise Storage Manager Topology Reporter (included in the Sun StorEdge Enterprise Storage Manager software) to be installed on a Solaris host. Refer to the “Other Supported Software” section in the release notes for your array for more information.

Note – This feature is not supported on a Linux host; that is, if the agent has been installed on a Linux host, the Configure Web Server option on the Custom Configure window is unavailable. You can use a Linux host web browser to load the console, but the agent must be installed on a system running the Solaris OS.

Web Browser Requirements

Sun StorEdge Enterprise Storage Manager Topology Reporter can be viewed through a URL on any machine connected to the same network (you must have an account on the Solaris host machine).

UNIX OS

- Netscape Navigator 4.79
- Java Plug-in 1.2.2_02

Note – For the IBM AIX OS, Java Plug-in software versions earlier than 1.3 are not supported.

Microsoft Windows OS

- Microsoft Internet Explorer (Versions earlier than 5.0 are not supported)
- Java Plug-in 1.2.2_02 software

or

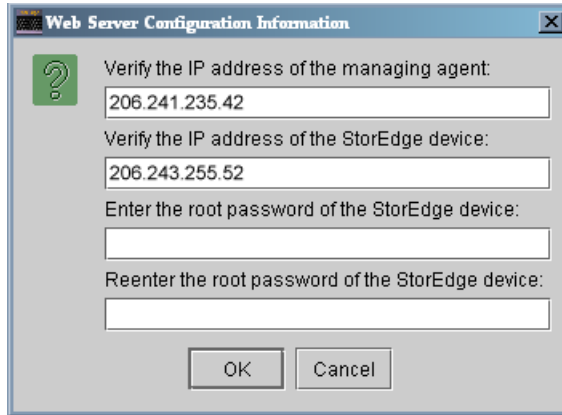
- Netscape Navigator 4.76
- Java Plug-in 1.2.2_02 software

Setting Up the Array

Although the entire Sun StorEdge Configuration Service package isn't required to be loaded on the machine that is going to be used to access the array from the web, a few program files are needed; therefore, the entire Sun StorEdge Configuration Service package must be installed on another server so you can transfer files from it to the array. The following procedure describes how to conveniently transfer the necessary files to set up web management.

1. **Make sure you have an Ethernet connection from the network to the array and you have established an IP address for it.**
2. **From the server that has the entire Sun StorEdge Configuration Service package installed, choose Configuration → Custom Configure.**
3. **Select Configure Web Server.**
4. **Verify the IP address of the managing agent, which is the agent that is running on the server directly connected to the storage.**
5. **Enter the IP address of the array where the files are going to be transferred and click OK.**

Note – By default there is no password set for the array. If you have created one using the firmware application, you must enter it. For information on creating or changing the password, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array.



After the files are transferred successfully, a confirmation message is displayed, and the console can now be accessed through the web.

▼ To Access the Console From the Web Browser

1. Open the web browser.
2. Type the following URL address:

`http://ip address of controller/esm.html`

3. Continue monitoring, maintaining, and updating storage as explained in the applicable chapters in this guide.


Maintaining the Array

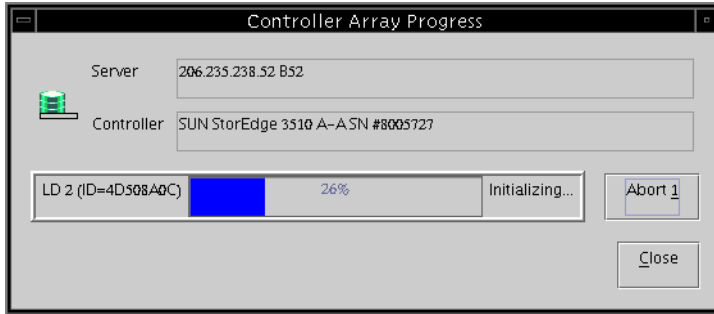
This chapter explains how to maintain the integrity of the array using Array Administration. Topics covered in this chapter are as follows:

- “Array Administration Activities” on page 111
 - “To Check Parity” on page 112
 - “To Schedule a Parity Check” on page 113
- “Failed Drives” on page 115
 - “To Automatically Rebuild a Drive Using a Standby Drive” on page 115
 - “To Rebuild a Device Without a Standby Drive” on page 116
 - “To Check the Progress of the Rebuilding Process” on page 117
 - “To Manually Rebuild a Failed Drive” on page 117
 - “To Restore a Logical Drive Configuration” on page 119
 - “To Reset the Controller” on page 122
 - “To Mute the Controller Beeper” on page 122
 - “To Bring a Failed Controller Back Online” on page 123
 - “To Convert a Dual Controller Array to a Single Controller Array” on page 123
 - “To Display Performance Statistics” on page 124

Array Administration Activities

Array administration activities such as initialization, drive rebuilding, and parity checking can take some time, depending on the size of the logical drive or physical drives involved.

After one of these processes has started, the Controller Array Progress window is displayed. If you close the window, to view progress, click the Progress Indicator icon  or choose View → Array Admin Progress.



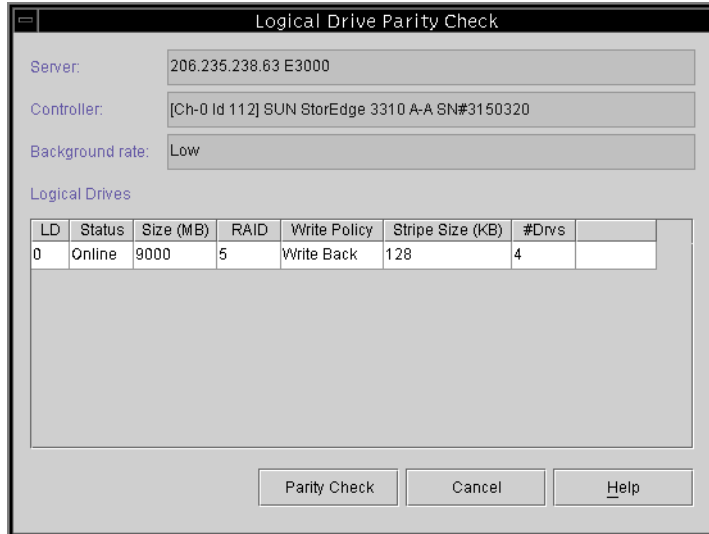
▼ To Check Parity

Parity checking is the process where the integrity of redundant data on fault-tolerant logical drives or LUNs (RAID 3 and 5) is checked.

1. **Select the logical drive on which you want to run parity check.**
2. **Choose Array Administration → Parity Check.**

Note – You need to be logged in as either `ssadmin` or `ssconfig` to access options on this menu.

3. **When the Logical Drive Parity Check window is displayed, select the logical drive on which you want to run a parity check.**



To run a parity check on multiple drives, choose the Array Administration → Schedule Parity Check command to schedule a parity check to be run in the near future (such as within three minutes). When scheduled parity check runs, it automatically performs the parity checks one after another.

4. Click the Parity Check button to start the parity check process.

Once a parity check has started, the Progress Indicator is automatically displayed. If this window is closed, it can be reopened by choosing View → Array Admin Progress or by clicking the Progress Indicator icon. A window is displayed that shows the percentage of completion progress for each array.

▼ To Schedule a Parity Check

Choose Array Administration → Schedule Parity Check to check parity of a specific logical drive array at scheduled intervals (for example, during off hours).

- You can choose to schedule any number of logical drives configured on a managed server; however, you can establish only one schedule per array controller.
- When you schedule multiple logical drives, the check is done in sequence from the lowest to the highest numbered logical drive.
- A parity check requires an average of five minutes per Gbyte of storage.
- You can control the amount of system resources allocated to the parity check by changing the rebuild priority. To change the rebuild priority, see [“Disk Array Tab” on page 160](#).

- A parity check can be stopped once it has started. There is no corruption as a result of this operation.

Note – You need to be logged in as either `ssadmin` or `ssconfig` to access options on this menu.

1. Select the controller on which you want to schedule the parity check.
2. Choose **Array Administration** → **Schedule Parity Check**.

The Schedule Parity Check window is displayed.

LD	Status	Size	RAID	Write Policy	Stripe Size	# of Drives
0	Online	9000	5	Write Back	128 KB	4

3. Make selections in the appropriate fields on this window.

- Listed Logical Drives – A list of the available fault-tolerant logical arrays. Use the Shift key to select multiple drives.
- Rebuild Priority – The amount of resources allocated to perform a rebuild (low, normal, improved, high). Select Change Controller Parameters from the Custom Configuration Options window to change the Rebuild Priority.
- How Often – Specify how often you want the parity checked.

- Starting Day – Specify the day of the week you would like this schedule to start.
- Starting Time – Specify the time on the starting day you would like this schedule to begin.
- Existing Schedule – The current parity schedule: logical drive, frequency, starting day, and time.
- Next Parity Check – The date and time you want the next parity check to start.

4. When you are satisfied with the schedule, click OK.

Failed Drives

This section contains procedures for recovering from a drive failure with and without a standby drive. If for some reason these procedures do not start the rebuilding process, instructions are also provided for manually starting a rebuild after a drive failure.



Caution – Install either local or global standby drives previously installed and configured for each logical drive before a failure occurs. Depending on the type of RAID level used and archiving procedure implemented, significant data loss might occur in cases of single or multiple drive failures. Additionally, make tested spare drives readily available on site for immediate replacement if a malfunction occurs.

▼ To Automatically Rebuild a Drive Using a Standby Drive

When a drive associated with a fault-tolerant logical drive fails, and a standby drive has previously been installed and configured as either a global or local spare, the failed drive is automatically substituted and its data rebuilt using the designated spare drive. For this to occur flawlessly, the spare drive's capacity must *always* be equivalent to or larger than the failed drive that is being replaced.

The rebuilding process normally starts within one to two minutes. It is performed in the background and takes approximately eight minutes per Gbyte when there is no other activity on the controller.

During the automatic rebuild process, normal activity might continue, although performance might degrade. The degree to which performance degrades is determined by the rebuild priority set for the controller. (To change the rebuild priority, see [“Disk Array Tab” on page 160.](#))

The progress of the rebuild process is displayed when you choose View → Array Admin Progress.

1. **Reestablish automatic rebuild capability by replacing the failed drive, using instructions contained in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.**

2. **Wait at least 60 seconds after removing the failed drive before inserting a new drive.**

Make sure the replacement drive is at least equal to the largest drive in the enclosure. Install the replacement drive in the same slot (drive bay) as the failed drive; the replacement drive then becomes the new standby drive.

3. **After the rebuild process is complete and the logical drive is online again, back up the array controller configuration to a file on an external drive or diskette.**

See [“To Save the Logical Drive Configuration” on page 27.](#)

▼ To Rebuild a Device Without a Standby Drive

If there is no standby drive in the array, you need to replace the failed drive before the automatic rebuild process can start.

1. **To recover from a drive failure when there is no standby drive, replace the failed drive by using the instructions contained in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.**

2. **Wait at least 60 seconds after removing the failed drive before inserting a new drive.**


Make sure the capacity of the replacement drive is at least equal to that of the failed drive. Install the replacement drive at the same address (drive bay) as the failed drive.

3. **Once the failed drive is replaced in the same slot, you need to scan it in.**

For detailed instructions on scanning in a drive, see [“To Scan in New Hard Drives” on page 146.](#)

4. **After the drive has been scanned, you need to manually rebuild it by choosing Array Administration → Rebuild.**

▼ To Check the Progress of the Rebuilding Process

1. **Choose View → Array Admin Progress or click the Progress Indicator icon in the upper right corner of the window .**

The Controller Array Progress window is displayed that shows the completion percentage of the rebuild. However, if there are activities (such as initialization, rebuild, or parity check) occurring on multiple controllers, the Select Controller Progress window is displayed first.

2. **Select the controller whose progress you want to view and click OK.**

The Controller Array Progress window is displayed that shows the array progress of the selected controller. For more information, see [“Array Administration Activities” on page 111](#).

▼ To Manually Rebuild a Failed Drive

In most cases, you do not need to use the manual rebuild process because replaced drives are automatically rebuilt.

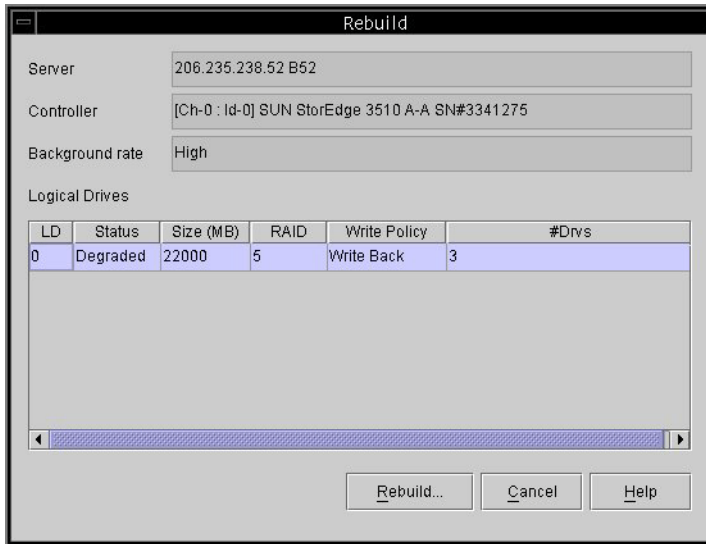
If a spare is not present when the failure occurs, or for some reason the drive does not rebuild, you can use Rebuild to manually start the rebuild process. Also, if the rebuild process is interrupted by a reset, use Rebuild to restart the rebuilding process.

1. **Replace the failed drive, using the instructions contained in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.**
2. **Wait at least 60 seconds after removing the failed drive before inserting a new drive.**

Make sure the capacity of the replacement drive is at least equal to that of the largest drive in the enclosure.

3. Choose Array Administration → Rebuild.

The Rebuild window is displayed.



4. Select the status record of the replacement drive.

5. Click Rebuild to start the rebuild process.

The rebuild process is performed in the background and takes approximately eight minutes per Gbyte when there is no other activity on the array controller. During a rebuild, normal activity can continue, although performance might degrade. The degree to which performance degrades is determined by the rebuild priority set for the controller. (To change the rebuild priority, see [“Disk Array Tab” on page 160.](#))

6. To check the progress of the rebuilding process, choose View → Array Admin Progress or click the Progress Indicator icon in the upper right corner of the window.

The Controller Array Progress window is displayed that shows the completion percentage of the rebuild.

If there are array activities (such as initialization, rebuild, or parity check) occurring on multiple controllers, the Select Controller Progress window is displayed first.

7. Select the controller whose progress you want to view and click OK.

The Controller Array Progress window is displayed and shows the array rebuilding status for that controller.

▼ To Restore a Logical Drive Configuration

This section describes how to restore the array configuration information from a backup file. You must have saved a backup file using the Save command as explained in [“Configuration File” on page 56](#). If the array controller and its drives are damaged, you can restore the array configuration to a new controller without having to completely reconfigure the storage array.

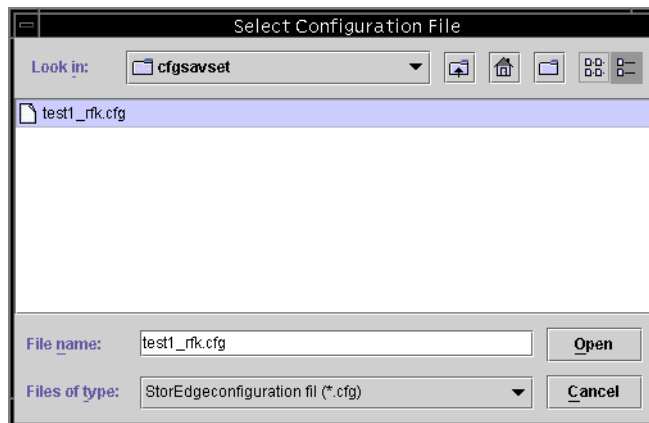


Caution – Restore the array configuration from a file only if the configuration file is current. Data loss will result from restoring an outdated or incorrect configuration.

If you are sure that the backup file contains the correct array configuration information, continue with the following procedure to restore the configuration.

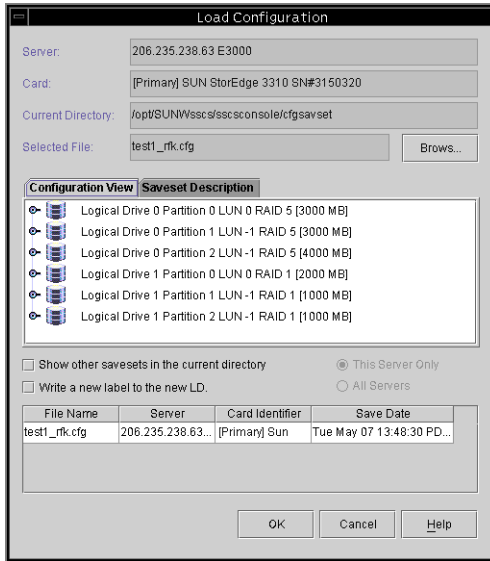
1. **Select the controller for the appropriate array.**
2. **Choose Configuration → Load Configuration.**

The Select Configuration File window is displayed.

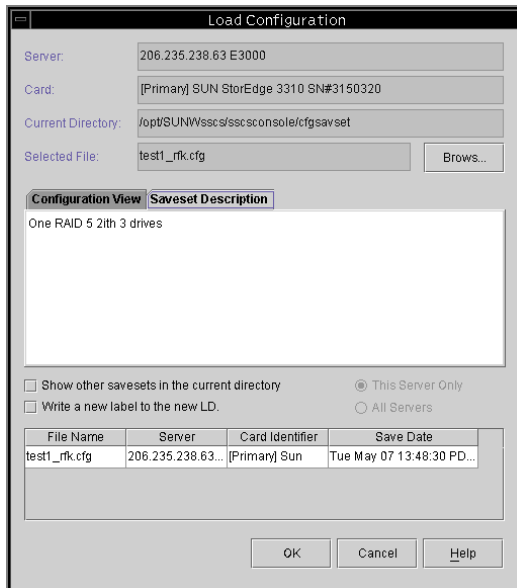


3. **Specify the name and location of the backup configuration file and click Open.**

The Load Configuration window is displayed. To see a tree-view representation of the configuration, click the Configuration View tab.



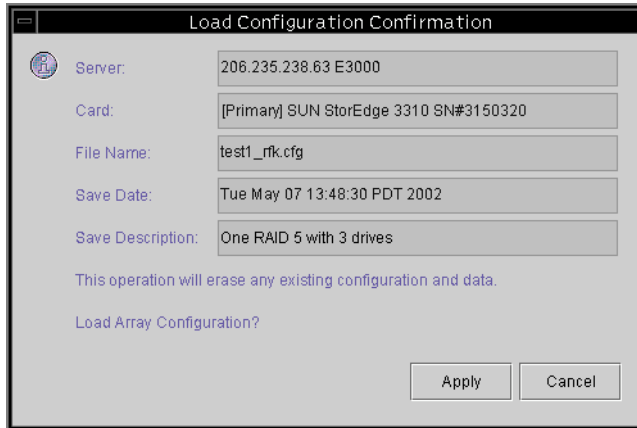
The Saveset Description tab displays the description of the file that was specified when the configuration file was created.



4. (Solaris OS only). For the SCSI and FC arrays, if you want the logical drive(s) to be automatically labeled, which enables the OS to use the drive, click **Write a new label to the new LD**.
5. To load the saved configuration, select **OK**.

The Load Configuration Confirmation window is displayed.

Carefully review the information presented in the Load Configuration Confirmation window before making a decision to continue.



6. Click **Apply** to load this configuration or click **Cancel** to terminate this function.

Apply causes the configuration operation to continue, and a progress window is displayed.

Note – Do *not* initialize LUN(s) after restoring the array configuration backup file contents.

7. If you are running Windows NT, shut down and restart the OS.

Reinitializing the OS is required to enable the logical drive configuration to be recognized by Windows NT.

▼ To Reset the Controller

Whenever you make changes to the controller parameters, you are asked if you want to reset the controller so that the changes take effect. If you are making multiple changes, you might not want to stop and reset the controller after each change. Use the Issue Reset to the Controller option to manually reset the controller after making multiple parameter changes.

1. **Select any storage icon in the main window.**
2. **Choose Array Administration → Controller Maintenance.**
3. **If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.**
The Controller Maintenance Options window is displayed.
4. **Click Issue Reset to the Controller.**

Note – Resetting the controller on a Sun StorEdge 3310 SCSI array can result in host-side error messages, such as parity error and synchronous error messages. No action is required and the condition corrects itself as soon as reinitialization of the controller is complete.

▼ To Mute the Controller Beeper

When an event occurs that causes the controller to beep, for example, when a logical drive fails, during a rebuild, or when adding a SCSI drive, you can mute the controller beeper. in one of two ways.

1. **Select the controller icon in the main window.**
2. **Choose Array Administration → Controller Maintenance.**
3. **If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.**
The Controller Maintenance Options window is displayed.
4. **Click Mute Controller Beeper.**
or
 1. **Select the desired controller icon in the main window.**
 2. **Choose Configuration → Custom Configure.**

3. **Select Change Controller Parameters.**
4. **Select Mute Beeper.**

Note – If the alarm is caused by a failed component, muting the beeper has no effect. You need to push the Reset button on the right ear of the array. See [“View Enclosure” on page 88](#) for more information about component failure alarms.

▼ To Bring a Failed Controller Back Online

If a controller fails, bring it back online in one of two ways.

1. **Select the desired controller icon in the main window.**
2. **Choose Array Administration → Controller Maintenance.**
3. **If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.**
The Controller Maintenance Options window is displayed.
4. **Click Deassert Failed Redundant Controller.**
or
 1. **Select the desired controller icon in the main window.**
 2. **Choose Configuration → Custom Configure.**
 3. **Select Change Controller Parameters.**
 4. **Select the Redundancy tab.**
 5. **From the Set Controller Config field, select Redundant Deassert Reset.**

▼ To Convert a Dual Controller Array to a Single Controller Array

If one controller fails in a dual array controller configuration, you might want to run a single controller for an extended period of time so that the array does not display as degraded in the console.

1. **Make sure you know the serial number of the controller being removed.**
You can check the event log for the failed controller’s serial number or check the console and make a note of the primary controller’s serial number.

2. Change the remaining controller's redundancy setting to disabled.

You must use the firmware application to disable redundancy on the controller. Refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array for information about accessing the firmware application, and then from the Main Menu, choose "view and edit Peripheral devices → Set Peripheral Device Entry → Redundant Controller – Primary → Disable redundant controller."

3. Stop the agent as explained in the installation chapter for your OS.

4. Change to `/var/opt/SUNWsscs/ssagent` and edit the file `sscontlr.txt`.

The very last line in the file contains the serial numbers of both controllers. Remove the failed controller's serial number from this line.

```
# RAID_CONTROLLER=Enable:3197861:3179746
```

5. Start the agent as explained in the installation chapter for your OS.

6. Rescan the console if it was open during this procedure.

▼ To Display Performance Statistics

Using Performance Statistics, you can determine the data transfer rate, that is, the speed the array is running at.

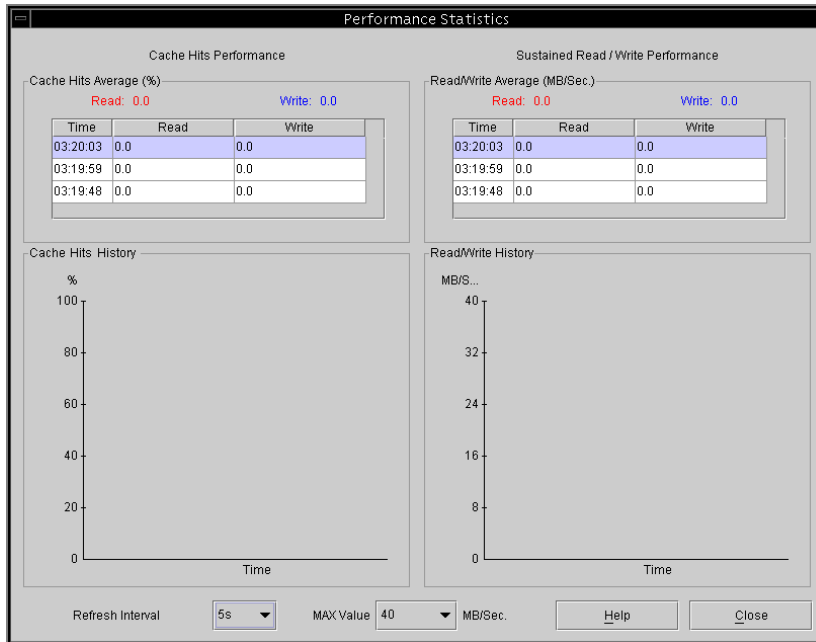
1. Choose Array Administration → Controller Maintenance.

2. If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.

The Controller Maintenance Options window is displayed.

3. Click Performance Statistics.


The Performance Statistics window is displayed.



Updating the Configuration

Refer to this chapter when you want to change the current configuration or add to it. This chapter describes the following tasks:

- “To Add a Logical Drive or Logical Volume From New Logical Drives” on page 128
- “To Add a Logical Volume From Existing Logical Drives” on page 132
- “To Delete a Logical Drive or Logical Volume” on page 133
- “To Create a Partition” on page 136
- “To Delete a Partition” on page 139
- “To Expand the Capacity of a Logical Drive or Logical Volume” on page 141
- “To Add SCSI Drives to an Existing Logical Drive” on page 143
- “To Copy and Replace Member Drives” on page 144
- “To Scan in New Hard Drives” on page 146
- “To Download RAID Controller Firmware” on page 147
- “To Upgrade Firmware and Boot Record” on page 149
- “To Upgrade Firmware on Hard Drives” on page 150
- “To Upgrade Firmware on SAF-TE/SES Devices” on page 152
- “To Change Controller Parameters” on page 153
- “To Save Changed Values” on page 154
- “To Mute the Controller Beeper” on page 167
- “To Assign or Change Standby Drives” on page 167
- “To Edit a Server Entry” on page 169
- “To Update the ODM” on page 171

The Configuration menu commands and tool icons might be temporarily disabled if an administration process, such as parity checking, is running. The menu command is also shown as deactivated when the console is refreshing its inventory on the server. A satellite dish symbol is attached to the server icon during the refresh process .

To use the Configuration options, you must log into the `ssconfig` security level of the software with the `ssconfig` password. When you are finished with the configuration activities, log back into the monitoring level of the program.

▼ To Add a Logical Drive or Logical Volume From New Logical Drives

Use this option to add one or more logical drives to an existing configuration of RAID sets, or for SCSI and FC arrays, to add a logical volume from *new* logical drives. For SCSI and FC arrays, to add a logical volume from *existing* logical drives, see [“To Add a Logical Volume From Existing Logical Drives” on page 132](#).

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).

Note – For SCSI and FC arrays, if the logical drive is going to be larger than 253 Gbyte, see [“To Prepare for Logical Drives Larger Than 253 Gbyte” on page 44](#).



Caution – On UNIX systems, before adding new logical drives to existing configurations, you must unmount any file systems that are attached to the array.

1. **Select the appropriate array.**
2. **Choose Configuration → Custom Configure.**

Note – This selection is inactive unless you have selected an array with available physical drives.

3. **Select Add LDs/LVs to the Current Configuration from the Custom Configuration Options window.**
4. **Verify that the server and controller displayed at the top of the window are correct.**
5. **Select a disk you want to be included in the new logical drive and click Add Disk.**
If you make a mistake or change your mind, select the drive and click Remove Disk.
6. **Select a RAID Level.**
For definitions of RAID levels, see [“RAID Basics” on page 173](#).
7. **Select the desired Host channel and SCSI ID to which you would like the new logical drive to be mapped to from the Channel and the SCSI ID list boxes.**

8. Set the Max Drive Size.

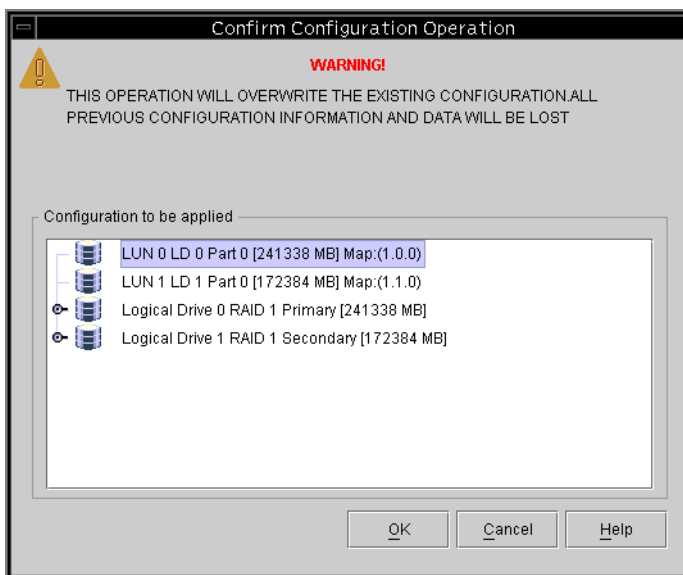
The Max Drive Size displays the total capacity of each disk. A smaller logical drive can be created by decreasing this value.

Note – If you do not change the Max Drive Size but you do change the Partition Size, a new partition is created at the specified partition size. The remaining logical drive size capacity moves to the last partition. Remaining capacity can be used later by expanding the drive (as explained in [“To Expand the Capacity of a Logical Drive or Logical Volume” on page 141](#)). The drive capacity is no longer editable after a partition is created.

Note – If you want to create another logical drive on the same controller, click New LD. The logical drive you just defined is created and you are returned to the top of the window, enabling you to create another logical drive. For the maximum number of logical drives supported, see [TABLE 4-1](#).

9. For SCSI and FC arrays, to add this logical drive to a logical volume, click New LD and see [“To Add a Logical Drive to a Logical Volume” on page 130](#).
10. When you are satisfied with the selections on this window, and do not want to define another logical drive, click Commit.

A confirmation window is displayed showing the new configuration.



11. Click OK to accept the configuration or Cancel to return to the console.

Note – You cannot change a logical drive configuration after you click OK.

Note – During initialization LD/LV size is displayed as 0 Mbyte.

12. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

13. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).

▼ To Add a Logical Drive to a Logical Volume

A logical volume is composed of two or more logical drives and can be divided into a maximum of 32 partitions for SCSI arrays and 128 for FC arrays. During operation, the host sees a nonpartitioned logical volume or a partition of a logical volume as one single physical drive.

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).



Caution – On UNIX systems, before adding new logical drives to existing configurations, you must unmount any file systems that are attached to the array.

1. Create a logical drive as described in Steps 1-8 in [“To Add a Logical Drive or Logical Volume From New Logical Drives” on page 128](#).

Note – Do not partition the logical drive that you are adding to the logical volume. A logical drive that has been partitioned cannot be added to a logical volume.

2. Before you click Commit, to add the logical drive to a logical volume, click Add to LV.

The logical drive is added to the LV Definition box. The total size of the logical volume is displayed in the Available Size (MB) field.

Note – Because the logical volume has not been partitioned yet, the Part Size (MB) and the Available Size (MB) are equal. A single logical volume is considered to be a single partition.

3. **To create another logical drive to add to the logical volume, click New LD.**
4. **Create the logical drive and add it to the logical volume by clicking Add to LV.**
Repeat this step for every logical drive you want to add to the logical volume.
5. **To create a partition, see [“To Create a Partition” on page 136](#).**
6. **When you have finished adding logical drives to the logical volume, to create another logical volume or an individual logical drive, click Commit LV.**
When you are finished creating logical volumes and do not want to create an individual logical drive, click Commit.

Note – When you are finished creating logical volumes and want to exit the New Configuration window, if you accidentally click Commit LV instead of Commit, you will have to create another logical drive; otherwise, you have to click Cancel and configure the logical volume again.

7. **(HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.**

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

8. **(IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).**

▼ To Add a Logical Volume From Existing Logical Drives

Note – Logical volumes are unsuited to some modern configurations such as Sun Cluster environments, and do not work in those configurations. Use logical drives instead. For more information, see [“Logical Volumes” on page 36](#).

Note – Before you can add a logical volume from existing logical drives, you must unmap the logical drives.

1. **Select the appropriate array.**
2. **Choose Configuration → Custom Configure.**

Note – This selection is inactive unless you have selected an array with available physical drives.

3. **Select Add LDs/LVs to the Current Configuration from the Custom Configuration Options window.**
4. **Verify that the server and controller displayed at the top of the window are correct.**
5. **Select Use existing LDs to create LVs.**

If you do not see any logical drives listed under Select a Logical Drive, the logical drives have not been unmapped and therefore are unavailable to select. You must unmap the logical drives first.

6. **Select a logical drive and click Add to LV.**
7. **When you have finished adding logical drives to the logical volume, to create another logical volume or an individual logical drive, click Commit LV.**

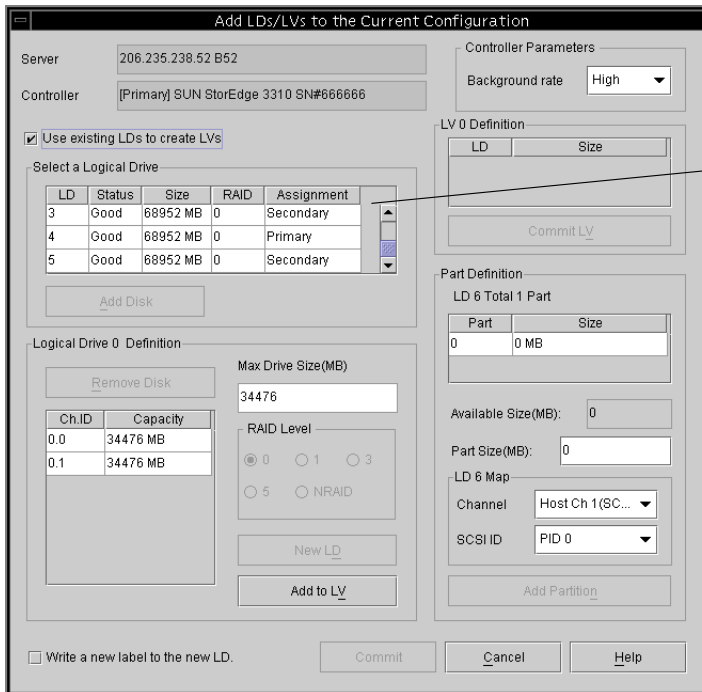
When you have finished creating logical volumes and do not want to create an individual logical drive, click Commit.

Note – When you are finished creating logical volumes and want to exit the New Configuration window, if you accidentally click Commit LV instead of Commit, you will have to create another logical drive; otherwise, you have to click Cancel and configure the logical volume again.

8. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

9. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).



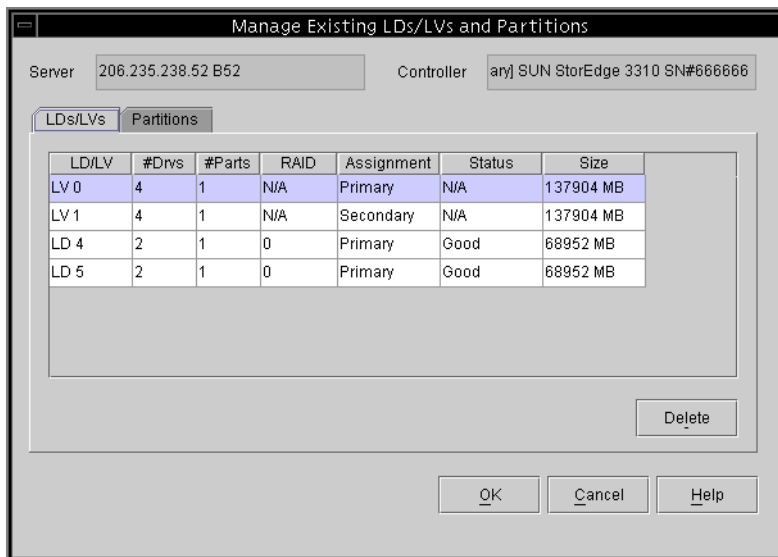
▼ To Delete a Logical Drive or Logical Volume

Use this option to delete one or more logical drives, or for SCSI and FC arrays, to delete logical volumes from an existing configuration of RAID sets.

Note – For SCSI and FC arrays, before you can delete a logical drive or logical volume, you need to unmap all assigned LUNs.

1. Select the appropriate array.
2. View the logical drives or logical volumes you want to delete.
3. If any of the logical drives or logical volumes have host LUN assignments, proceed to [Step 4](#); if they do not, proceed to [Step 8](#).
4. Choose Configuration → Custom Configure.
5. Select Change Host LUN Assignments.
6. Select the host LUNs attached to the logical drive or logical volume you want to unmap, and click Unmap Host LUN.
7. Click OK, and then click Close.
8. Choose Configuration → Custom Configure.
9. Select Manage Existing LDs/LVs and Partitions from the Custom Configuration Options window.
10. Select the LDs/LVs tab.
11. Select the logical drive or logical volume you want to delete, and click Delete and click OK.

When deleting a logical volume, after you click Delete, the logical volume is deleted, but the logical drives that make up the logical drive are displayed.



12. Click OK in the Confirm Configuration Operation window to complete the operation, and click Close.

13. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

14. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171](#).

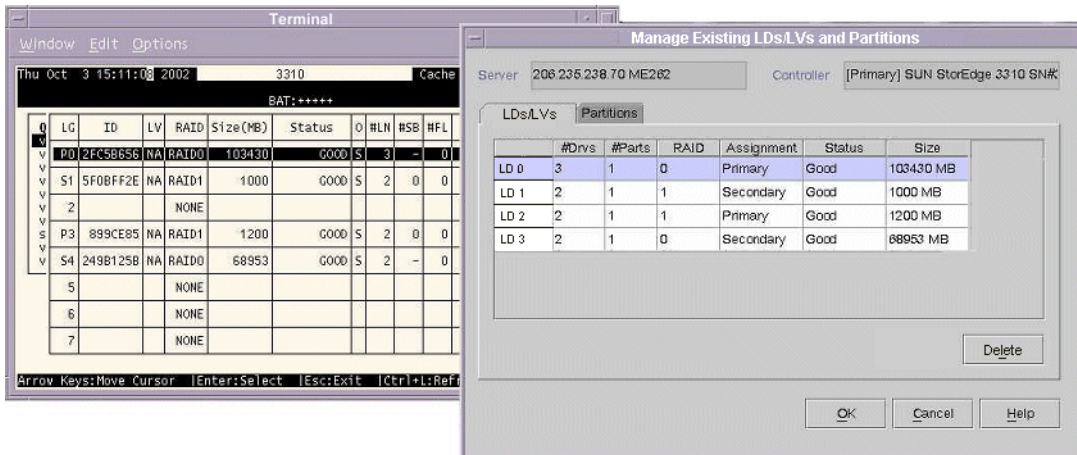
The Logical Drive/Logical Volume Number

The logical drive/logical volume number referenced with each logical drive is dynamic; it changes when logical drives are created/deleted. This number is displayed in the logical drive (LDs/LVs) field of several windows including Dynamically Grow and/or Reconfigure LDs/LVs, Change Host LUN Assignments, Manage Existing LDs/LVs and Partitions, and the main window.

Used strictly as a placeholder that enables you to *visually* keep track of logical drives and logical volumes, this number is insignificant to the controller. That is, the controller does not report on the logical drives or logical volumes according to this number. For example, if four logical drives exist, and LD2 is deleted, the existing LD3 dynamically changes to LD2, and LD4 changes to LD3. Only the LD/LV number changes; all LUN mapping and data on the logical drives remains unchanged.

Because the controller reports on the *total number* of logical drives, which in this case is three, the actual LD/LV number as displayed in the LD/LV field is irrelevant. In this example, if a new logical drive is created, it takes the LD number of the logical drive that was deleted, and the controller reports that there are a total of four logical drives. All existing logical drives return to their original primary/secondary designation.

Note – As shown in the following example, the LG number on the firmware terminal menu option View and Edit Logical Drives is *not* visually dynamic. After a logical drive is deleted, you see an empty placeholder. When a logical drive is created from the console or from the terminal, this empty placeholder is filled with the new logical drive.



The LD field in the Manage Existing LDs/LVs and Partitions window (right) is dynamic; it changes when logical drives are created/deleted. In this example, LD2 was deleted and LD4 became LD3 and LD 2 became LD1. In the terminal menu option View and Edit Logical Drives (left), the LG field is not dynamic; an empty slot indicates that LD2 was deleted.

▼ To Create a Partition

Note – For SCSI and FC arrays, before you can create a partition, you need to unmap all assigned LUNs.

1. Select the appropriate array.
2. View the logical drive(s) you want to create partition(s) on.
3. If any of these logical drives have host LUN assignments, proceed to [Step 4](#); if they do not, proceed to [Step 8](#).
4. Choose Configuration → Custom Configure.
5. Select Change Host LUN Assignments.
6. Select the Host LUN(s) that are attached to the logical drive(s) you want to partition, and click Unmap Host LUN.
7. Click OK, and then click Close.
8. Choose Configuration → Custom Configure.

9. Select **Manage Existing LDs/LVs and Partitions from the Custom Configuration Options window.**

10. Select the **Partitions** tab.

11. Select a logical drive or logical volume you want to partition.

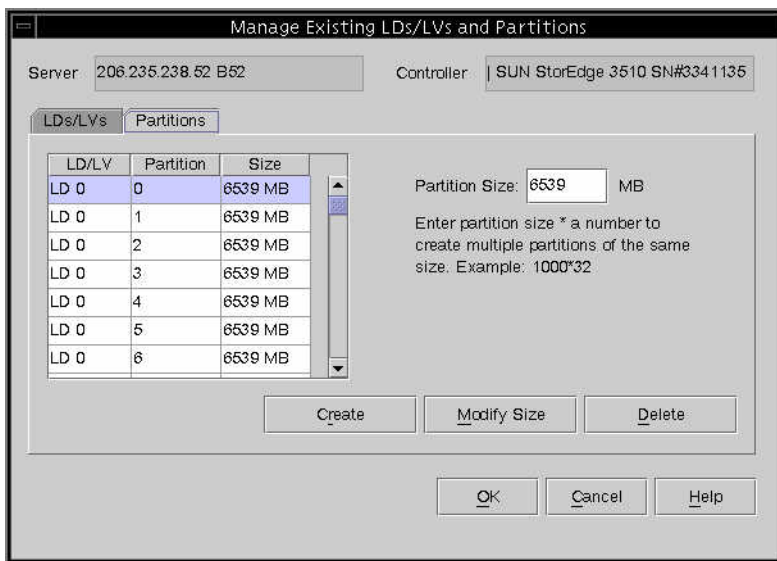
12. Specify the **Partition Size in Mbyte** and click **Create**.

To create multiple partitions of the same size, click **Add Partition** as many times as partitions you want to create. You can also type the partition size in the **Part Size** field and multiply (*) it by the number of partitions you want to create, for example 100*128. Any remaining Mbyte is added to the last partition.


As you add partitions, the remaining capacity displayed in **Available Size (MB)** decreases by the amount of the partition size.

13. To change the size of a partition you have already created, select the logical drive or logical volume, and click **Modify Size**.

14. Specify the new size (in Mbyte) in the **Partition Size** field, and click **OK**.




15. Click **OK** in the **Confirm Configuration Operation** window to complete the operation, and click **Close**.

After a logical drive or logical volume has been partitioned, when you open a logical drive or logical volume on the main window, the partitions are displayed .

16. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

17. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171.](#)



18. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

19. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host” on page 171.](#)

The Logical Drive/Logical Volume Number

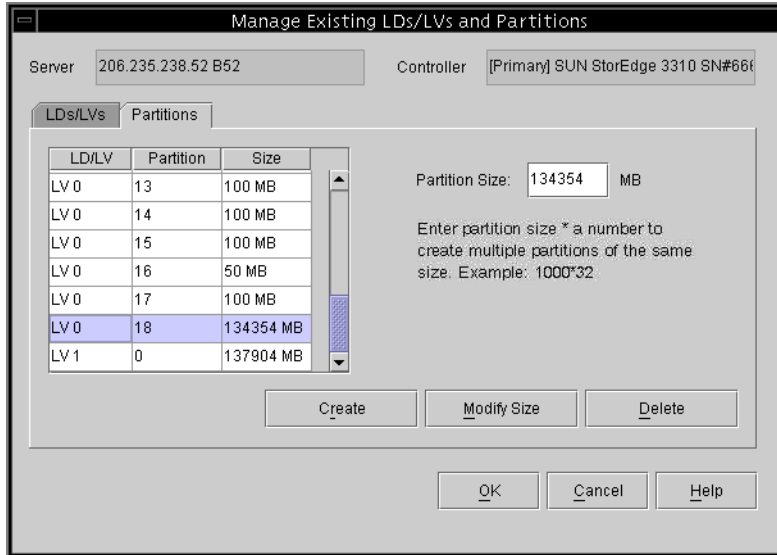
For important information regarding the logical drive/logical volume number displayed in the LDs/LVs field in the Manage Existing LDs/LVs and Partitions window, see [“The Logical Drive/Logical Volume Number” on page 135.](#)

▼ To Delete a Partition

Note – For SCSI and FC arrays, to delete a partition on a logical drive or logical volume, you need to unmap all assigned LUNs.

1. **Select the appropriate array.**
2. **View the logical drives or logical volumes on which you want to delete the partitions.**
If any of the partitions on the drive have host LUN mappings, proceed to [Step 3](#); if they do not, proceed to [Step 7](#).
3. **Choose Configuration → Custom Configure.**
4. **Select Change Host LUN Assignments.**
5. **Select the LUNs that are mapped to the logical drive's or logical volume's partitions that you want to delete, and click Unmap Host LUN.**
6. **Click OK, and then click Close.**
7. **Choose Configuration → Custom Configure.**
8. **Select Manage Existing LDs/LVs and Partitions from the Custom Configuration Options window.**
9. **Select the Partitions tab.**
10. **Select a partition to modify or delete by starting from the last partition within the logical drive or logical volume.**

11. Click Delete, and then click OK.



12. Click OK in the Confirm Configuration Operation window to complete the operation, and click Close.
13. (HP-UX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to run the `ioscan -fnC disk` command.

Note – If you used System Administrator Manager (sam) to unmount the file system, make sure it is closed before running the `ioscan` command.

14. (IBM AIX OS only). To ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM) as explained in [“Updating the Object Data Manager on an IBM AIX Host”](#) on page 171.

▼ To Expand the Capacity of a Logical Drive or Logical Volume

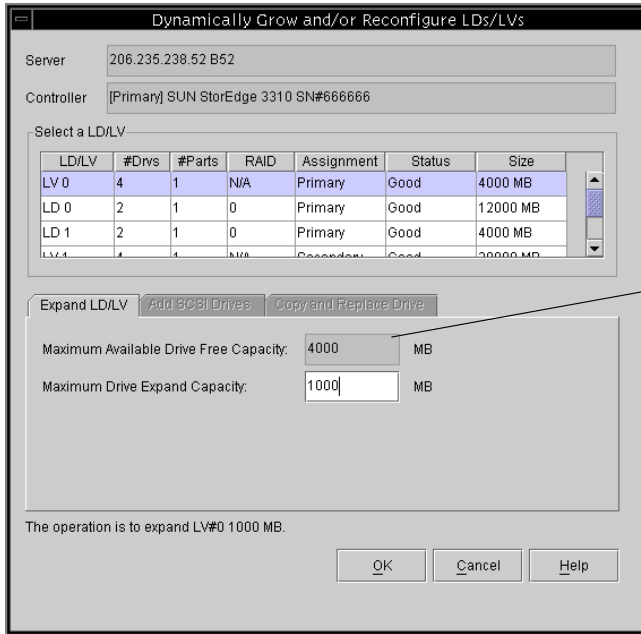
Use this option to expand the capacity of an existing logical drive, or for SCSI and FC arrays, to expand the capacity of a logical volume. For example, you might originally have had an 18-Gbyte drive of which only 9 Gbyte were selected to build a logical drive. To use the remaining 9 Gbyte, you need to expand the logical drive.

Note – For SCSI and FC arrays, to expand a logical volume, you must first expand the logical drives that make up the logical volume.

1. Select the appropriate array.
2. Choose Configuration → Custom Configure.
3. Select Dynamically Grow and/or Reconfigure LDs/LVs from the Custom Configuration Options window.
4. Select the logical drive or logical volume you want to expand.
5. Select the Expand LD/LV tab.

- Specify the capacity in Mbyte by which you want to expand the logical drive or logical volume in the Maximum Drive Expand Capacity field, and click OK.

The Maximum Drive Expand Capacity cannot exceed the Maximum Available Drive Free Capacity.



If you used the full capacity of the logical drive when it was created, 0 is displayed; there is no room to expand unless you add another logical drive.

- Click OK in the Confirm Configuration Operation window to complete the operation, and click Close.

The Logical Drive/Logical Volume Number

For important information regarding the logical drive/logical volume number displayed in the LD/LV field in the Dynamically Grow and/or Reconfigure LDs/LVs window, see [“The Logical Drive/Logical Volume Number”](#) on page 135.

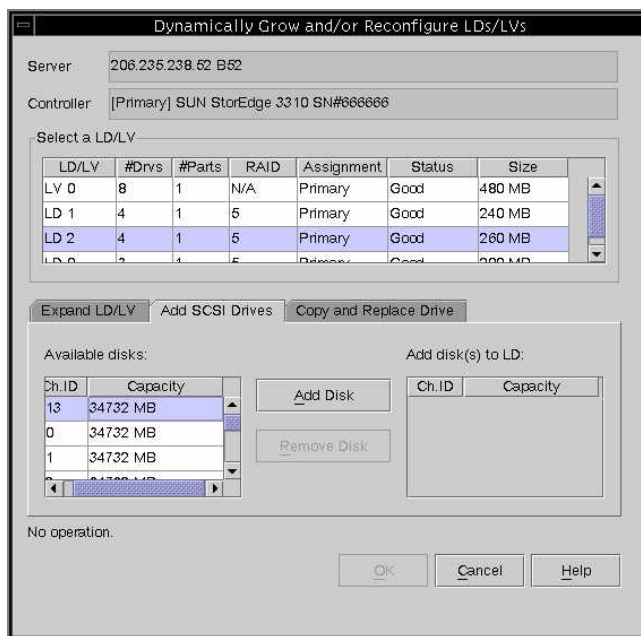
▼ To Add SCSI Drives to an Existing Logical Drive

1. Select the appropriate array.
2. Choose Configuration → Custom Configure.
3. Select Dynamically Grow and/or Reconfigure LDs/LVs from the Custom Configuration Options window.
4. Select the logical drive to which you want to add a SCSI drive.
5. Select the Add SCSI Drives tab.
6. From the list of Available disks, select the drive you want to add to the logical drive.
7. Click Add Disk.

The drive is moved to the Add disk(s) to LD list.

If you make a mistake or change your mind, select the disk from the Add disk(s) to LD list and click Remove.

8. When you are finished adding the SCSI drives, click OK.



9. Click OK in the Confirm Configuration Operation window to complete the operation, and click Close.

The Logical Drive/Logical Volume Number

For important information regarding the logical drive/logical volume number displayed in the LD/LV field in the Dynamically Grow and/or Reconfigure LDs/LVs window, see [“The Logical Drive/Logical Volume Number”](#) on page 135.

▼ To Copy and Replace Member Drives

You can copy and replace existing member drives with drives of the same or higher capacity. Because the logical drive uses the capacity size of its smallest drive, *all* drives must be replaced with drives of the same or higher capacity. For example, as shown [FIGURE 8-1](#), a logical drive that originally contains three 36-Gbyte member drives can be replaced with new 73-Gbyte member drives.

Note – To use the additional capacity provided by drives with higher capacity, you need to expand the capacity as explained in [“To Expand the Capacity of a Logical Drive or Logical Volume”](#) on page 141.

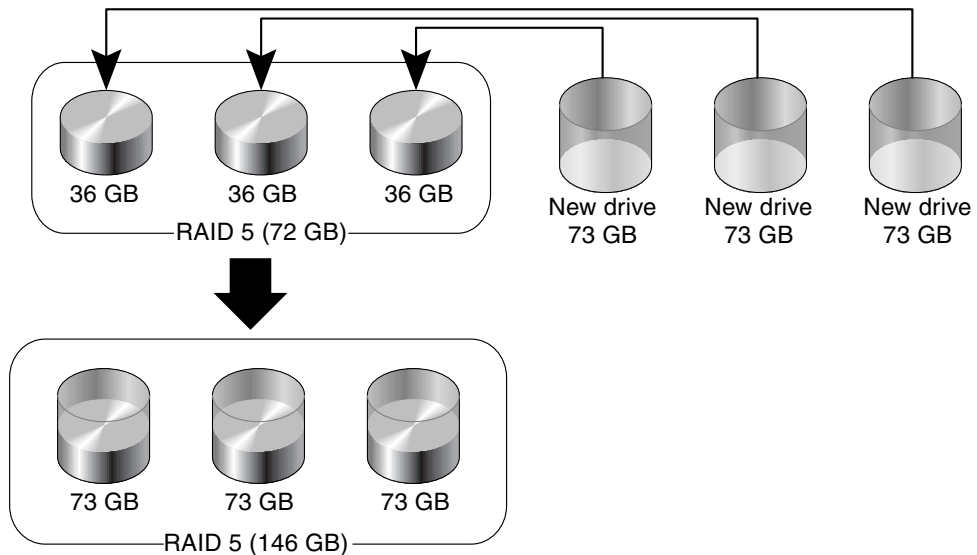
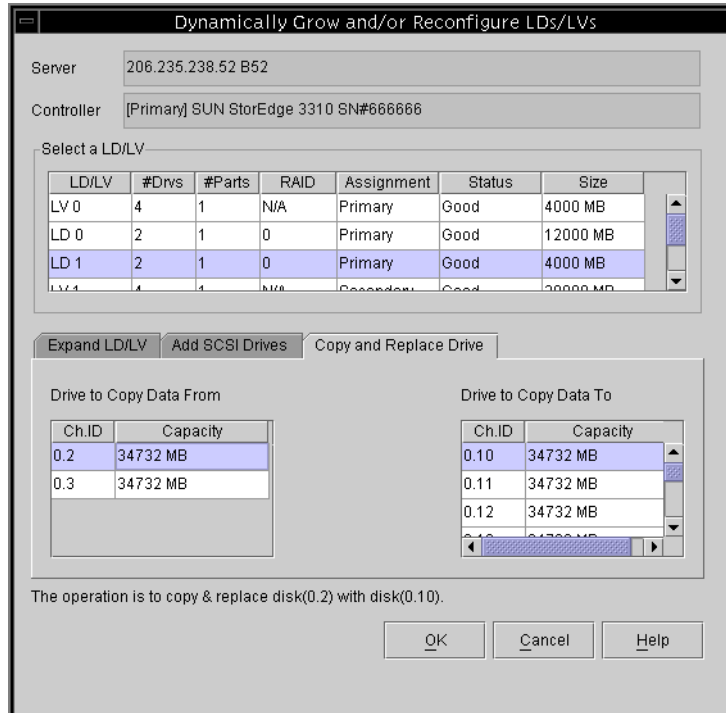


FIGURE 8-1 Copying and Replacing Member Drives

1. Select the appropriate array.
2. Choose Configuration → Custom Configure.
3. Select Dynamically Grow and/or Reconfigure LDs/LVs from the Custom Configuration Options window.
4. Select the logical drive for which you are going to perform the copy and replace operation.
5. Select the Copy and Replace Drive tab on the Dynamically Grow and/or Reconfigure LDs/LVs window.
6. From the Drive to Copy Data From list, select the new hard drive.
7. From the Drive to Copy Data To list, select the hard drive that is going to be replaced, and click OK.



8. Click OK in the Confirm Configuration Operation window to complete the operation, and click Close.
9. When the operation is complete, close the progress window.

10. To use the additional capacity provided by the new drives, follow the instructions in [“To Expand the Capacity of a Logical Drive or Logical Volume”](#) on page 141.

The Logical Drive/Logical Volume Number

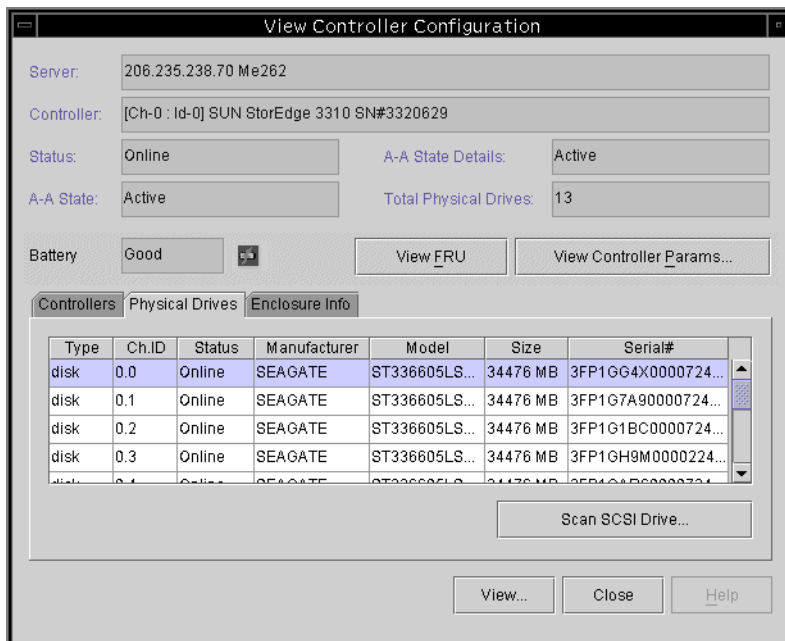
For important information regarding the logical drive/logical volume number displayed in the LD/LV field in the Dynamically Grow and/or Reconfigure LDs/LVs window, see [“The Logical Drive/Logical Volume Number”](#) on page 135.

▼ To Scan in New Hard Drives


A hard drive can be scanned in and made available without having to shut down the array.

1. Double-click the appropriate array.
2. The View Controller Configuration window is displayed.
3. Select the Physical Drives tab, and click Scan SCSI Drive.

If a drive fails, the Scan SCSI Drive button is also displayed on the Physical Drive window. You can select a physical drive, select View, and click Scan SCSI Drive from the View Physical Drive window.



4. Select the correct Channel and ID that the drive was input on.



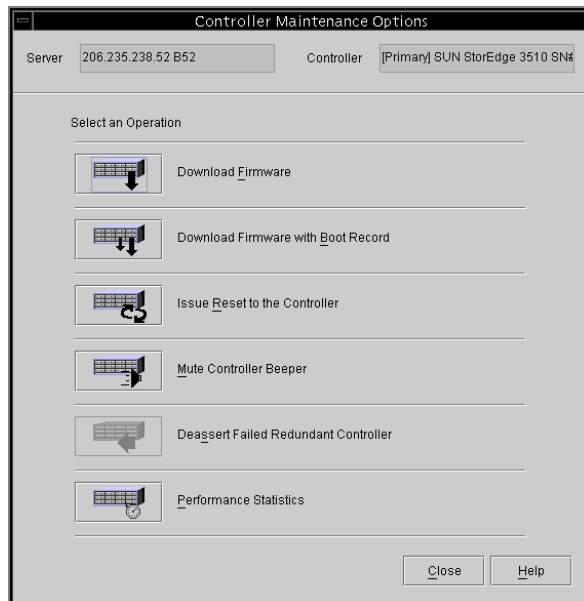
If the scan was successful, the drive is appropriately displayed and made available.

▼ To Download RAID Controller Firmware

The following procedures are used to upgrade the controller firmware for both a single and redundant controller configuration.


1. Select the desired controller icon in the main window.
2. Choose **Array Administration** → **Controller Maintenance**.
3. If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.

The Controller Maintenance Options window is displayed.



4. If upgrading firmware only (not boot record), select the Download Firmware option.

The Select Firmware File window is displayed.



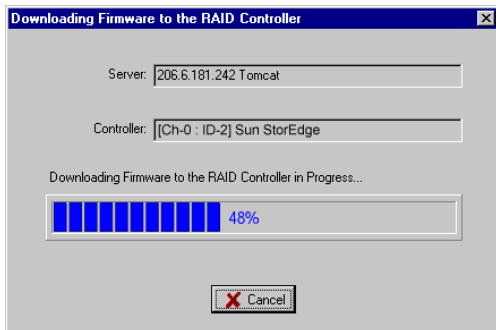
5. Select the firmware you want to download, and click Open.

The Confirmation Dialog prompt is displayed.



6. Click Yes.

The firmware download to the RAID controller displays a progress bar.



7. When the progress bar reaches 100%, click OK.

8. After the firmware has been downloaded, check the settings to make sure they are configured correctly.

▼ To Upgrade Firmware and Boot Record

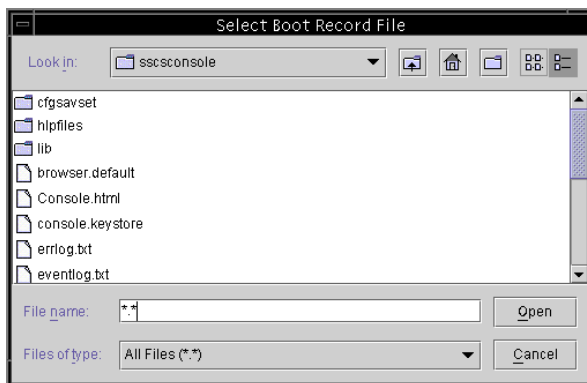
1. Choose Array Administration → Controller Maintenance.

2. If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.

The Controller Maintenance Options window is displayed.

3. Select Download Firmware with Boot Record.

The Select Boot Record File window is displayed.



4. **Select the boot record and click Open.**
5. **Select the appropriate firmware file.**
The Select Firmware File is displayed.
6. **Click Open.**
The Confirmation Dialog window is displayed.
7. **Repeat Steps 6 through 8 in the previous subsection.**

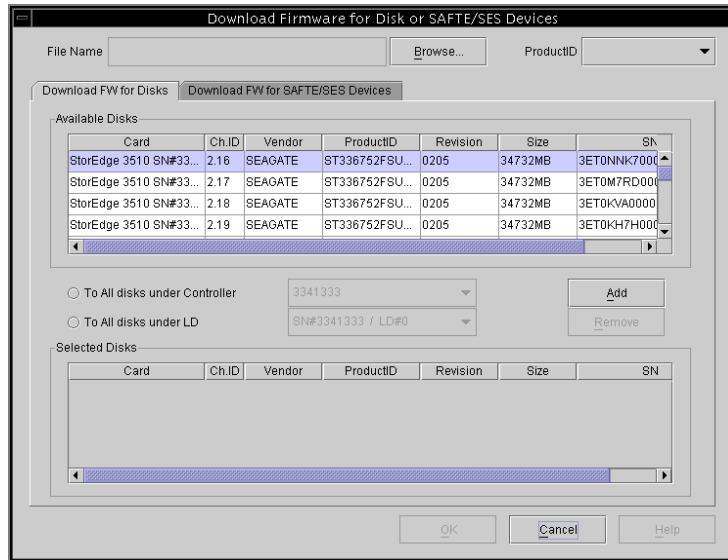
Downloading Firmware for Devices

This option enables you to upgrade the firmware on hard drives and SAF-TE/SES devices.

▼ To Upgrade Firmware on Hard Drives

1. **Select the appropriate array.**
2. **Choose Array Administration → Download FW for Devices.**
3. **Click the Download FW for Disks tab.**
4. **Select either To All disks under Controller, and select an array from the menu, or select To All disks under LD, and select the logical drive from the menu.**
 - If there are drives that you do not want the new firmware downloaded to, select them from Selected Disks, and click Remove.
 - If there are logical drives that you want to add, select them from Available Disks and click Add.

- If you have multiple drives that have different product IDs, make sure you select the product ID for the drive(s) to which you are downloading firmware from the Product ID list box.

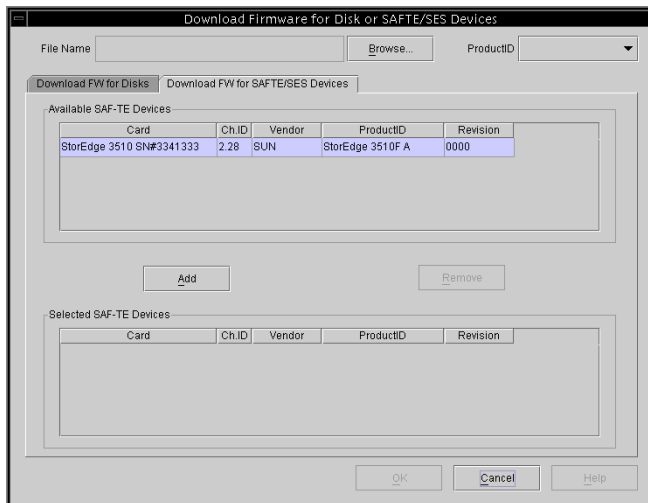


5. Click **Browse** and locate the download firmware file.
Select **Open**.
6. Select the download firmware file, click **Open**, and click **OK**.
The firmware starts to download.
7. When the progress reaches 100%, click **OK**.
8. To verify that the firmware has downloaded successfully, select **View** → **View Physical Drive**, and make sure the firmware version has changed in the **Product Revision** field.

▼ To Upgrade Firmware on SAF-TE/SES Devices

Note – SAF-TE devices are used by SCSI arrays and SES devices are used by Fibre Channel arrays.

1. Select the appropriate array.
2. Choose **Array Administration** → **Download FW for Devices**.
3. Click the **Download FW for SAFTE/SES Devices** tab.
 - To add a device, select it from Available SAF-TE Devices and click Add.
 - To remove a device, select it from Selected SAF-TE Devices and click Remove.



4. Click **Browse** and locate the download firmware file.
5. Select the download firmware file, click **Open**, and click **OK**.
The firmware starts to download.
6. When the progress reaches 100%, click **OK**.
7. To verify that the firmware has downloaded successfully, select **View** → **View Enclosure**, and make sure the firmware version has changed in the **Firmware Rev** field.

▼ To Change Controller Parameters

1. Select the appropriate array.
2. Choose Configuration → Custom Configure.

If necessary, log in to the configuration level of the program with the `ssconfig` password. The Custom Configuration Options window is displayed.

3. From the Custom Configuration Options window, select Change Controller Parameters.

The Change Controller Parameters window with the Channel tab selected is displayed.

Chl #	LogChl...	Mode	Type	PID	SID	DefClk	DefWid	Term	CurClk	CurWid
0	0	Drive	SCSI	6	7	80 MHz	Wide	Off	80 MHz	Wide
1	0	Host	SCSI	0	1	80 MHz	Wide	Off	80 MHz	Wide
2	1	Drive	SCSI	6	7	80 MHz	Wide	Off	Async	Narrow
3	1	Host	SCSI	3	2	80 MHz	Wide	Off	80 MHz	Wide
6	0	N/A	RCCOM	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Note – For the Sun StorEdge 3510 FC array, the CurClk is 2.0 GHz.

- Controller Name – If the controller name needs to be set, select Controller Name and type the desired name. Click OK to save the change.
- Controller Unique ID – This ID is automatically set.

▼ To Save Changed Values

Some of the options on the Change Controller Parameters window require that the controller be reset so that the changes take effect. If a change requires a controller reset, the following message is displayed in the lower left side of the window:

[Controller reset is required for changes to take effect.]

To reset the controller and save changed values, you can either select the Issue Controller Reset check box at the time of making the change, or reset the controller later through the Controller Maintenance window (see [“To Reset the Controller” on page 122.](#)). If you are making multiple changes, you might not want to stop and reset the controller after each change. If you do not select the Issue Controller Reset check box, and the change requires a reset, when you click OK, a warning message is displayed:

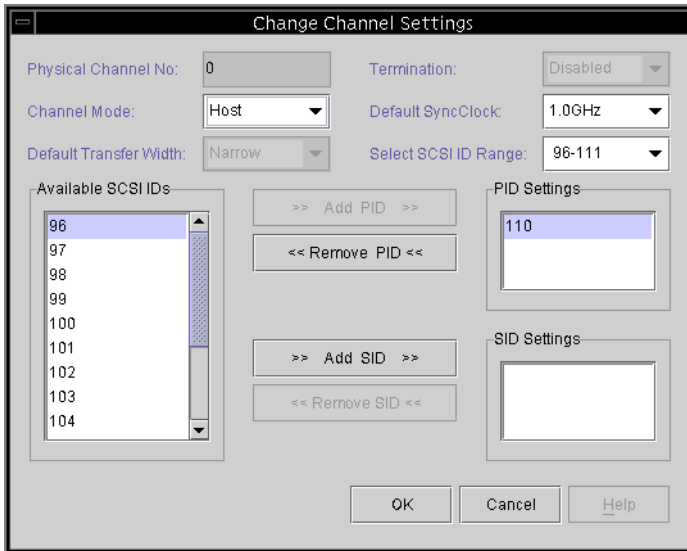


1. Select the Issue Controller Reset check box.
 2. Make the changes and click OK.
- or
1. Do not select the Issue Controller Reset check box.
 2. Make the changes and click OK.
 3. Reset the controller later as explained in [“To Reset the Controller” on page 122.](#)

Channel Tab

1. From the Channel Settings tab, select the channel to be edited.
2. Click Change Settings.

The Change Channel Settings window is displayed. For the server to recognize the array, a host channel must have an ID assigned to a logical drive and a logical drive mapped to that host channel and ID. This window enables you to configure the host/drive channel.



3. From the Channel Mode list box, select either Host or Drive.

A Drive channel is what the drives are connected to (internal or external). A host channel is what is connected to the server.

Note – The Sun StorEdge 3310 SCSI array supports a maximum of two host channels.

Note – Depending on the controller configuration, you might need to select both primary and secondary channel IDs as described in the following steps.

4. From the Available SCSI IDs list box, select the primary channel ID, which is designated as PID, and click Add PID.
5. If you have two controllers installed, select a secondary channel ID from the Available SCSI IDs list box, and click Add SID.

Note – For the Sun StorEdge 3310 SCSI array, If you add more than four host channel IDs, the LUNs Per Host ID parameter (see “Host I/F Tab” on page 163) must be set to a value less than 32.

6. For changes to take effect, reset the controller.

Changing Host ID in a Fibre Configuration

1. If you want an ID higher than 15, click Select SCSI ID Range and select the desired range.

Note – Each channel’s ID must be within the same range.

2. Click Remove to remove the PID or SID.
3. Once your selections have been made, click OK to redisplay the previous window.

RS 232 Tab

1. After all channel settings have been made, from the Change Controller Parameters window, select the RS 232 tab.

The screenshot shows the 'Change Controller Parameters' window with the 'RS 232' tab selected. The window title is 'Change Controller Parameters'. At the top, there are fields for 'Server: 206.235.238.65 SUNblade' and 'Controller: [Primary] SUN StorEdge 3310'. Below this is the 'Controller Data' section with various fields: 'Firmware Version: 3.230', 'Boot Record Version: 1.21F', 'Serial Number: 3197861', 'CPU Type: PPC750', 'Controller Name: 7100', 'Cache Size: 2MB ECC SDRAM', 'Controller Unique ID: 6', and checkboxes for 'Mute Beeper' and 'Restore Factory Defaults'. There is also an 'Issue Controller Reset' checkbox. Below the controller data is a tabbed interface with 'RS 232' selected. Underneath is a table with the following data:


Port #	Max Baud Rate	Min Baud Rate	Default Baud Rate	Current baud Rate
0	38400	2400	9600	38400
1	38400	2400	9600	38400

At the bottom of the window are 'Change Settings...', 'OK', and 'Cancel' buttons.

2. Select the port desired, and click Change Settings.

The Change RS232 Port Settings window is displayed.

3. Select any baud rate desired including the default setting of 9600, and then click OK to return to the previous window.



Cache Tab

Optimization

From the Cache tab, you can select the optimization mode for the logical drive, which indicates the amount of data that is written across each drive in an array. Sequential I/O indicates large blocks of data and Random I/O indicates small blocks of data. [TABLE 8-1](#) specifies the stripe size for each RAID level, depending on the optimization mode.

TABLE 8-1 Stripe Size Per Optimization Mode

RAID Level	Sequential I/O	Random I/O
0, 1, 5	128	32
3	16	4

The type of application the array is working with determines whether Random or Sequential I/O should be applied. Video/imaging application I/O size can be 128, 256, 512 Kbyte, or up to 1 Mbyte, so the application reads and writes data to and from the drive as large-block, sequential files. Database/transaction-processing applications read and write data from the drive as small-block, randomly-accessed files.

There are two limitations that apply to the optimization modes:

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

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

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

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

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

TABLE 8-2 Maximum Number of Disks per Logical Drive for a 2U Array

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

TABLE 8-3 Maximum Usable Capacity (Gbyte) per Logical Drive for a 2U Array

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

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

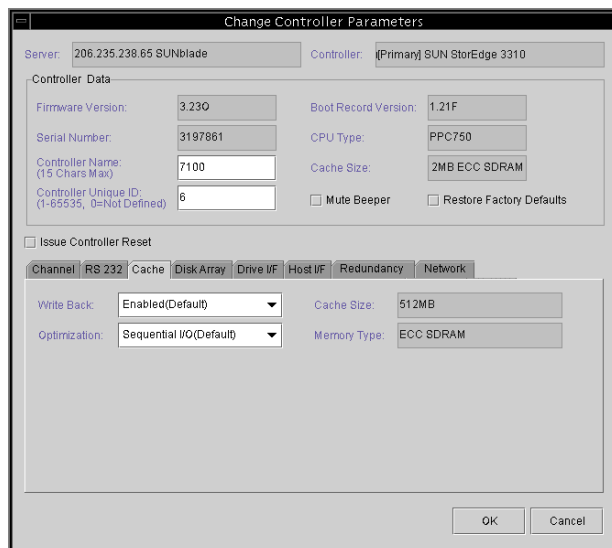
Write-Back Cache

Write-back is a cache-writing strategy in which the controller receives the data to be written to disk, stores it in the memory buffer, and immediately sends the host OS a signal that the write operation is complete, without waiting until the data is actually written to the disk drive. Write-back caching improves the performance of write operations and the throughput of the controller card. Write-through is a cache-writing strategy in which the controller writes the data to the disk drive before signaling the host OS that the process is complete. Write-through has lower write operation and is considered more secure if power failure should occur. Because a battery module is installed, power is supplied to the data cached in memory and the cached writes can be completed when power is restored.

1. From the Change Controller Parameters window, select the Cache tab.
2. Select Optimization from the list box or accept the current setting.
3. To specify Write Back cache, click the list box and select Enabled.

If Write Back is disabled, Write-Through is the selected cache-writing strategy.

4. For changes to take effect, reset the controller.



Disk Array Tab

1. **From the Change Controller Parameters window, select the Disk Array tab.**
2. **Select either Disabled or Enabled from the three Write Verify list boxes.**

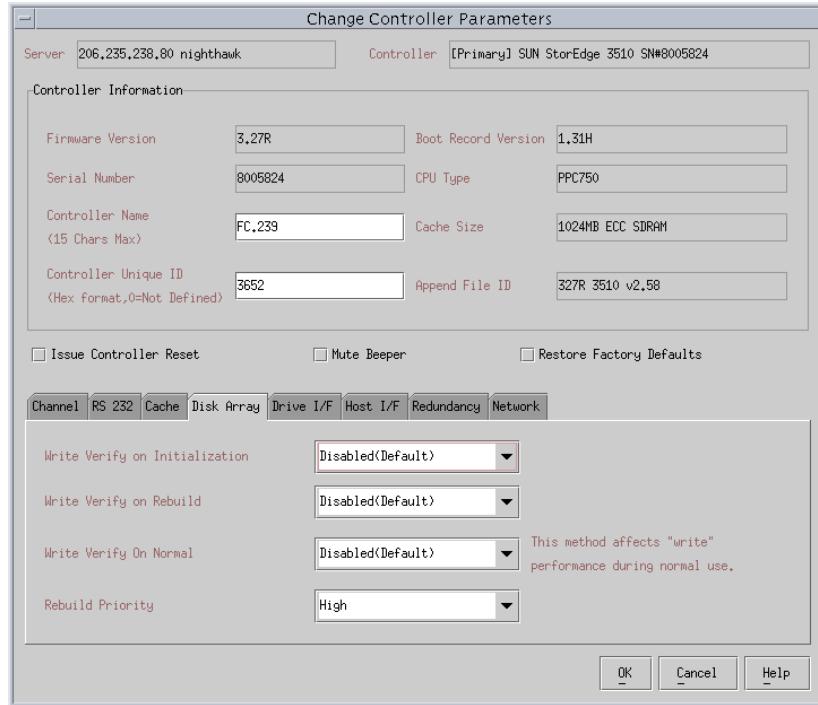
Normally, errors might occur when a hard drive writes data. To avoid the write error, the controller can force the hard drives to verify the written data.

- Write Verify on Initialization - Performs verify-after-write while initializing the logical drive.
- Write Verify on Rebuild – Performs verify-after-write during the rebuilding process.
- Write Verify on Normal – Performs verify-after-write during normal I/O requests.

3. **Select from the four options available in the Rebuild Priority list box: Low, Normal, Improved, or High.**

The RAID controller provides a background rebuilding ability. This means the controller is able to serve other I/O requests while rebuilding the logical drives. The time required to rebuild a drive set largely depends on the total capacity of the logical drive being rebuilt. Additionally, the rebuilding process is totally transparent to the host computer or the OS.

- Low – The default that uses the controller's minimum resources to rebuild.
- Normal – To speed up the rebuilding process.
- Improved – To allocate more resources to the rebuilding process.
- High – to use the controller's maximum resources to complete the rebuilding process in the shortest time possible.



Drive I/F Tab

1. From the Change Controller Parameters window, select the Drive I/F tab.
2. From the SCSI Motor Spin Up field, select either Disabled or Enabled.

The SCSI Motor Spin Up decides how the SCSI drives in a Disk Array are started. When the power supply is unable to provide sufficient current for the hard drives and controllers that are powered up at the same time, spinning up the hard drives serially is one of the best ways to consume lower powerup current. By default, all hard drives spin-up when powered on.

3. From the Power Up SCSI Reset field, select either Enabled or Disabled.

By default, when the controller is powered up, it sends a SCSI bus reset command to the SCSI bus. When disabled, it does not send a SCSI bus reset command on the next power up.

When connecting dual host computers to the same SCSI bus, the SCSI bus reset interrupts all the read/write requests being performed. This might cause an OS or host computers to act abnormally. Disable Power Up SCSI Reset at power up to avoid this situation.

4. Set the Disk Access Latency.

This function sets the delay time before the controller tries to access the hard drives after power on. The default is 15 seconds.

5. Set the Tag Count Per drive.

This is the maximum number of tags that can be sent to each drive at the same time. A drive has a built-in cache that is used to sort all of the I/O requests (tags) that are sent to the drive, enabling the drive to finish the requests faster.

The cache size and maximum number of tags varies between different brands and models of drive. Use the default setting of 32. Changing the maximum tag count to Disable causes the internal cache of the drive to be ignored (not used).

The controller supports tag command queuing with an adjustable tag count from 1 to 128.

6. From the SAF-TE Polling Period (Sec) field, select the variable time options shown in the list box, or select 0.0 to disable this function so that all installed Event Recording Modules (ERMs) are never polled.

7. From the SCSI I/O Timeout (Sec) field, select from 0.5 through 30 seconds.

The SCSI I/O Timeout is the time interval for the controller to wait for a drive to respond. If the controller attempts to read data from or write data to a drive but the drive does not respond within the SCSI I/O timeout value, the drive is considered a failed drive. The default setting for SCSI I/O Timeout is 15 seconds for the SCSI array and 30 seconds for the Fibre Channel array.



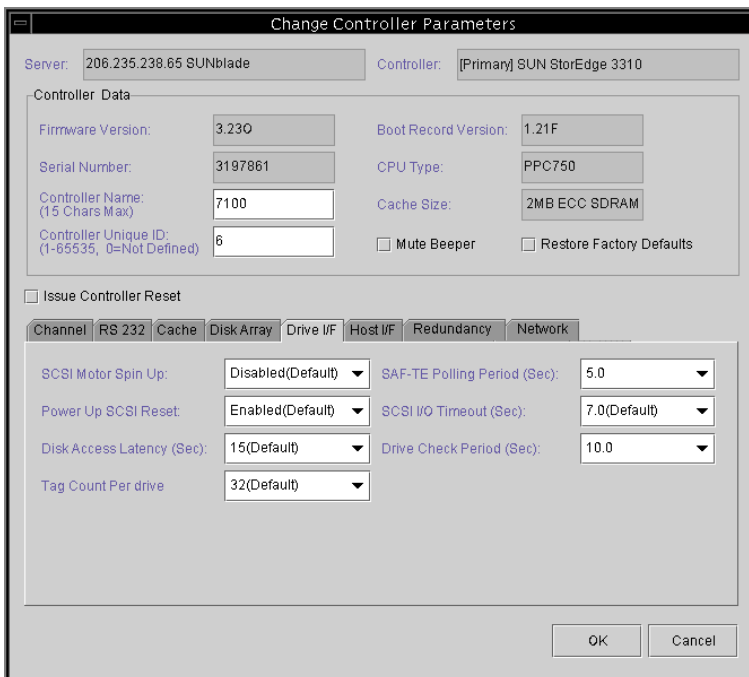
Caution – Do not change this setting. Setting the timeout to a lower value causes the controller to judge a drive as failed while a drive is still retrying or while a drive is unable to arbitrate the SCSI bus. Setting the timeout to a greater value causes the controller to keep waiting for a drive, and it might sometimes cause a host timeout.

When the drive detects a media error while reading from the drive platter, it retries the previous reading or recalibrates the head. When the drive encounters a bad block on the media, it reassigns the bad block to another spare block on the same disk drive. However, all of this takes time. The time to perform these operations can vary between different brands and models of drives.

During SCSI bus arbitration, a device with higher priority can use the bus first. A device with lower priority sometimes receives a SCSI I/O Timeout when devices of higher priority keep using the bus.

8. From the Drive Check Period (Sec) field, select from 0.5 through 30 seconds.

The Periodic Drive Check Time is an interval for the controller to check the drives on the SCSI bus. The default value is Disabled, which means if there is no activity on the bus, the controller does not know if a drive has failed or has been removed. Setting an interval enables the program to detect a drive failure when there is no array activity; however, performance is degraded.



Host I/F Tab

- 1. From the Change Controller Parameters window, select the Host I/F tab.**
- 2. Set the Max Queued IO Count.**

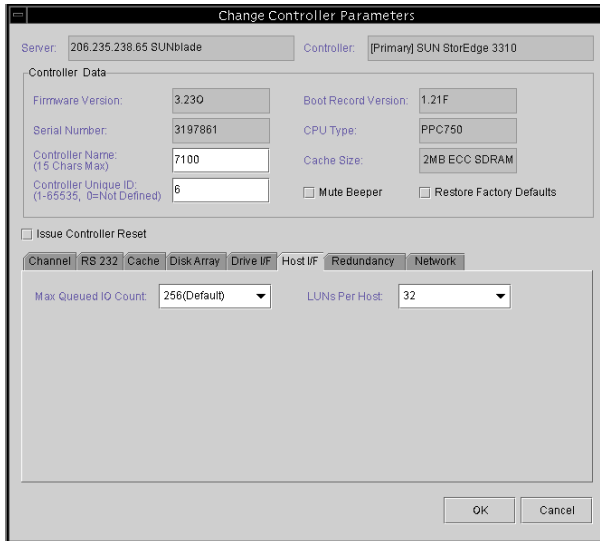
This function enables you to configure the maximum size of the I/O queue the controller can accept from the host computer by byte size. The default is 256. The predefined range is from 1 to 1024 bytes, or you might choose the Auto Computer (Automatically Configured) mode.

- 3. Set the LUNs Per Host.**

This function is used to change the number of LUNs per host SCSI ID. Each time a host channel ID is added, it uses the number of LUNs allocated in this setting. The default setting is 32 LUNs, with a predefined range of 1 to 32 LUNs available.

Note – For the Sun StorEdge 3310 SCSI array, the maximum number of LUN assignments is 128; therefore, if you use the default setting of 32, you can only add four host channel IDs ($4 \times 32 = 128$). If you added more than four host channel IDs (see [“Channel Tab” on page 155](#)), the LUNs Per Host parameter must be set to a value less than 32.

4. For changes to take effect, reset the controller.



Redundancy Tab

1. From the Change Controller Parameters window, select the Redundancy tab.

Three read-only fields are displayed: Controller Configuration, Controller Status, and Secondary Serial No.

2. Select an option from the Set Controller Config field.

- Redundant Deassert Reset – If you have failed a controller and want to bring it back online.
- Redundant Force Sec Fail – If you want to force the secondary controller to fail.
- Redundant Force Pri Fail – If you want to force the primary controller to fail.

The screenshot shows a 'Change Controller Parameters' dialog box. At the top, it displays 'Server: 206.235.238.65 SUNblade' and 'Controller: [Primary] SUN StorEdge 3310'. Below this is a 'Controller Data' section with fields for Firmware Version (3.23Q), Boot Record Version (1.21F), Serial Number (3197861), CPU Type (PPC750), Controller Name (7100), Cache Size (2MB ECC SDRAM), and Controller Unique ID (6). There are also checkboxes for 'Mute Beeper' and 'Restore Factory Defaults'. A section for 'Issue Controller Reset' is present but unchecked. At the bottom, there are tabs for 'Channel', 'RS 232', 'Cache', 'Disk Array', 'Drive I/F', 'Host I/F', 'Redundancy', and 'Network'. The 'Network' tab is active, showing 'Controller Configuration: Redundant Primary', 'Controller Status: Redundant Enabled', 'Secondary Serial No: 3179746', and 'Set Controller Config: No Change'. 'OK' and 'Cancel' buttons are at the bottom right.

Note – Set both controllers in the Redundant Primary configuration. The controllers then determine which one is primary and which one is secondary. This prevents any possible conflicts between controllers.

3. For changes to take effect, reset the controller.
4. Click OK to return to the main menu.

Network Tab

1. From the Change Controller Parameters window, select the Network tab.
2. To manually configure an IP address, subnet mask, or gateway address, click Change Settings.
3. If you have set up an array in an environment with a DHCP/RARP Server, you can select Enable Dynamic IP Assignment to enable the array to automatically obtain an IP address from the network.

The screenshot shows the 'Change Controller Parameters' window with the 'Network' tab selected. The window title is 'Change Controller Parameters'. At the top, there are two text boxes: 'Server: 206.235.238.70 Me262' and 'Controller: [Primary] SUN StorEdge 3310'. Below this is the 'Controller Data' section with several fields: 'Firmware Version: 3.25E', 'Boot Record Version: 1.31C', 'Serial Number: 3320629', 'CPU Type: PPC750', 'Controller Name: (15 Chars Max) [empty]', 'Cache Size: 2MB ECC SDRAM', 'Controller Unique ID: (1-65535, 0=Not Defined) Not Defined', and two checkboxes: 'Mute Beeper' and 'Restore Factory Defaults'. Below the Controller Data is a checkbox 'Issue Controller Reset'. At the bottom of the window are 'OK' and 'Cancel' buttons.

Channel RS 232 Cache Disk Array Drive I/F Host I/F Redundancy Network

Chl #	MAC Address	Static IP Address/NetMask/Gateway	Dynamic IP ...	Dynam...
7	00:C0:FF:32:AB:35	206.235.238.200 / 255.255.255.0 / 206.235.238.1	Disable	206.235.238.2

Change Settings...

OK Cancel

▼ To Mute the Controller Beeper

When an event occurs that causes the controller to beep, for example, when a logical drive fails, during a rebuild, or when adding a SCSI drive, you can mute the beeper in one of two ways.

1. **Select the desired controller icon in the main window.**
2. **Choose Array Administration → Controller Maintenance.**
3. **If you are not already logged in as `ssconfig`, a password prompt is displayed; type the `ssconfig` password.**
The Controller Maintenance Options window is displayed.
4. **Click Mute Controller Beeper.**
or
1. **Select the desired controller icon in the main window.**
2. **Choose Configuration → Custom Configure.**
3. **Select Change Controller Parameters.**
4. **Select Mute Beeper.**

Note – If the alarm is caused by a failed component, muting the beeper has no effect. You need to push the Reset button on the right ear of the array. See [“View Enclosure” on page 88](#) for more information about component failure alarms.

▼ To Assign or Change Standby Drives

A standby drive acts as a spare to support automatic data rebuilding after a physical drive in a fault-tolerant (non-RAID 0) logical drive fails. For a standby drive to take the place of another drive, it must be at least equal in size to the failed drive and all of the logical drives dependent on the failed disk must be redundant (RAID 1, 3, 5, or 1+0).

With this function you can either assign a global or local standby drive or change a ready drive’s state to standby or a standby drive’s state to ready. A drive that is assigned as a global spare rebuilds if a member of any existing drive fails. You can have one or more standby drives associated with an array controller. Global spares are used in the order in which they are created. A local spare has to be assigned to a particular logical drive and only rebuilds for a member within that logical drive.

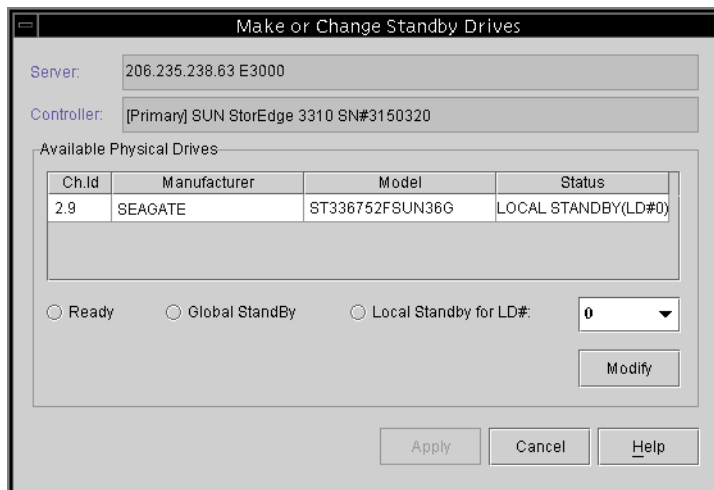
1. **In the main window, select the desired array controller.**

2. Choose Configuration → Custom Configure or click the Custom Configuration tool.

If necessary, log into the configuration level of the program with the `ssconfig` password. The Custom Configuration Options window is displayed.

3. Select Make or Change Standby Drives from the Custom Configuration Options window.

The Make or Change Standby Drives window is displayed.



4. Check the server and the controller IDs at the top of the window.

If you want to select a different server or controller, click Cancel to return to the main window, select the correct server, or controller from the tree view, and repeat Steps 2 and 3.

5. Select a drive to be assigned or changed.
6. Change or assign the drive's state by selecting Ready, Global StandBy, or StandBy for LD# (local).
7. Click Modify.
8. Click Apply, and then click Close.
9. Whenever you make changes to the configuration, save the new configuration to a file. For details, see ["Configuration File" on page 56](#).


Available Servers

Occasionally, you might need to edit or delete an entry from the Available or Managed Servers lists in the Server List Setup window.

▼ To Edit a Server Entry

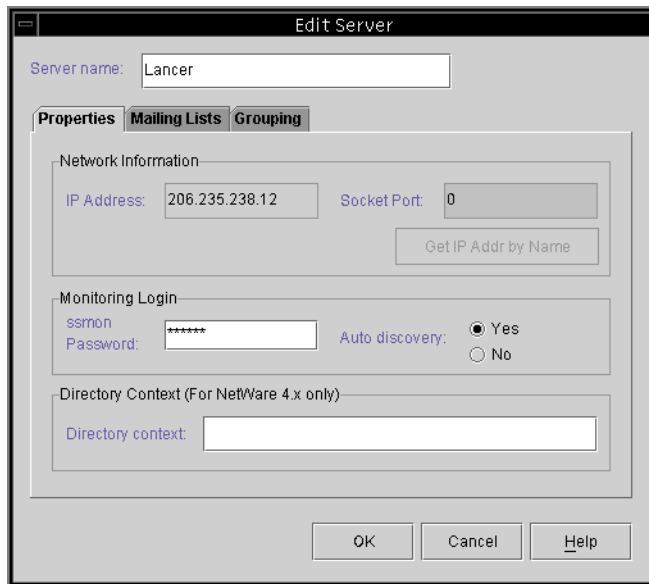
1. Choose **File** → **Server List Setup**. The **Server Setup** window is displayed.

If necessary, move the server name from the Managed Servers list to the Available Servers list in the Server List Setup window. Note that only the server entries in the Available Servers list can be edited.



2. Select the name of the server in the Available Servers list, and click Edit.

The Edit Server window is displayed.



3. Make the necessary changes. Click OK to register your changes.

For descriptions of the fields in this window, see [“To Add Servers” on page 12](#). The Add Server and Edit Server windows contain the same fields.

IP Address Shortcut: If the network address has changed, click Get IP Addr by Name. The program searches for and displays the correct IP address if you typed the name of the server as it is recorded by the name service used by the network.

If the name used for the server is not the same as the server’s network name or if the naming service is not yet updated, delete the server and add it again.

4. Move the server name back to the Managed Servers list.

5. Click OK to exit the Edit Server window.

Updating the Object Data Manager on an IBM AIX Host

For an IBM AIX host, to ensure the environment is stable and accurate after making configuration changes, you need to update the Object Data Manager (ODM).

▼ To Update the ODM

1. Run the following command for each deleted disk:

```
# rmdev -l hdisk# -d
```

where # is the number of the disk that was removed.



Caution – Never remove hdisk0.

To remove multiple disks (hdisk1 up to hdisk19), run the following commands:

```
# /usr/bin/ksh93
# for ((i=1; i<20; i++))
> do
> rmdev -l hdisk$i -d
> done
```

If the `rmdev` command returns disk busy errors, use either the command line, `smit`, or `smitty` to make sure that any previously created volume groups have been varied off and that no file systems are mounted on the device(s). It might also be necessary to perform an `exportvg` on persistent volume groups. If `exportvg` does not work, try rebooting.

2. If using a JBOD, run the same command for generic devices, which can be determined from the results returned from running the following command:

```
# lsdev -Cc generic
```

3. Run the following commands:

```
# /usr/bin/ksh93
# for ((i=1; i<20; i++))
> do
> rmdev -l gsc$i -d
> done
```

4. Delete references in the /dev directory by running the command:

```
# rm /dev/gsc*
```

5. Stop and start the agent and reread the system configuration into the ODM by running the following commands:

```
# ssagent stop
# ssagent start
# cfgmgr -v
```



Caution – Depending on the number of devices present in the OS, this command might take several minutes to complete. Do not make any configuration changes until `cfgmgr` has completed.

RAID Basics

This appendix provides background information about RAID including an overview of RAID terminology and RAID levels. Topics covered include the following:

- [“RAID Terminology Overview” on page 173](#)
- [“RAID Levels” on page 179](#)
- [“Local and Global Spare Drives” on page 185](#)

RAID Terminology Overview

Redundant array of independent disks (RAID) is a storage technology used to improve the processing capability of storage systems. This technology is designed to provide reliability in disk array systems and to take advantage of the performance gains offered by an array of multiple disks over single-disk storage.

RAID’s two primary underlying concepts are:

- distributing data over multiple hard drives improves performance
- using multiple drives properly allows for any one drive to fail without loss of data and without system downtime

In the event of a disk failure, disk access continues normally and the failure is transparent to the host system.

Logical Drive

A logical drive is an array of independent physical drives. Increased availability, capacity, and performance are achieved by creating logical drives. The logical drive appears to the host the same as a local hard disk drive does.




FIGURE A-1 Logical Drive Including Multiple Physical Drives

Logical Volume

A logical volume is composed of two or more logical drives. The logical volume can be divided into a maximum of 32 partitions for SCSI and 128 partitions for Fibre Channel. During operation, the host sees a nonpartitioned logical volume or a partition of a logical volume as one single physical drive.

Local Spare Drive

A local spare drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the local spare drive becomes a member drive and automatically starts to rebuild.

Global Spare Drive

A global spare drive does not only serve one specified logical drive. When a member drive from any of the logical drives fails, the global spare drive joins that logical drive and automatically starts to rebuild.

SCSI Channel

A SCSI channel (SCSI bus) can connect up to 15 devices (excluding the controller itself) when the Wide function is enabled (16-bit SCSI). Fibre Channel enables the connectivity of up to 125 devices in a loop. Each device has one unique SCSI ID.

FIGURE A-2 illustrates this idea. If you file documents into a cabinet, you must put the document into one of the drawers. From a SCSI point of view, a SCSI ID is like a cabinet, and the drawers are the LUNs (logical units). Each SCSI ID enables up to 32 LUNs. Data can be stored into one of the LUNs of the SCSI ID. Most SCSI host adapters treat a LUN like another SCSI device.




FIGURE A-2 SCSI ID/LUNs

A logical drive consists of a group of SCSI or Fibre Channel drives. Physical drives in one logical drive do not have to come from the same SCSI channel. Also, each logical drive can be configured for a different RAID level.

A drive can be assigned as the local spare drive to one specified logical drive, or as a global spare drive. A spare is not available for logical drives that have no data redundancy (RAID 0).

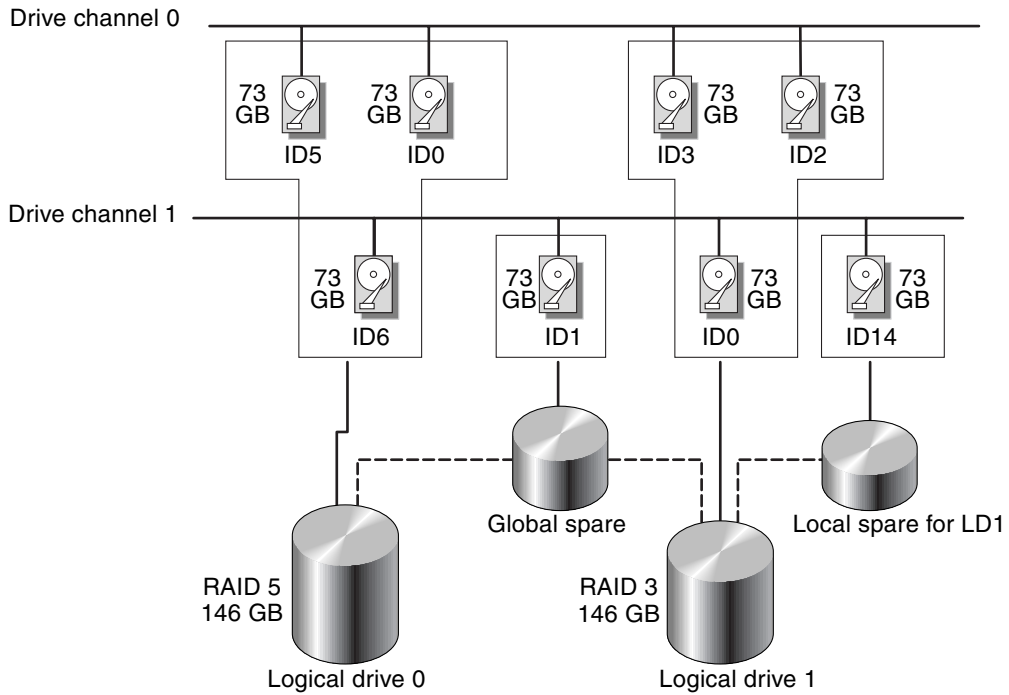


FIGURE A-3 Allocation of Drives in Logical Drive Configurations

You can divide a logical drive or logical volume into several partitions or use the entire logical drive as single partition.




FIGURE A-4 Partitions in Logical Drive Configurations

Each partition is mapped to LUNs under host SCSI IDs or IDs on host channels. Each SCSI ID/LUN acts as one individual hard drive to the host computer.

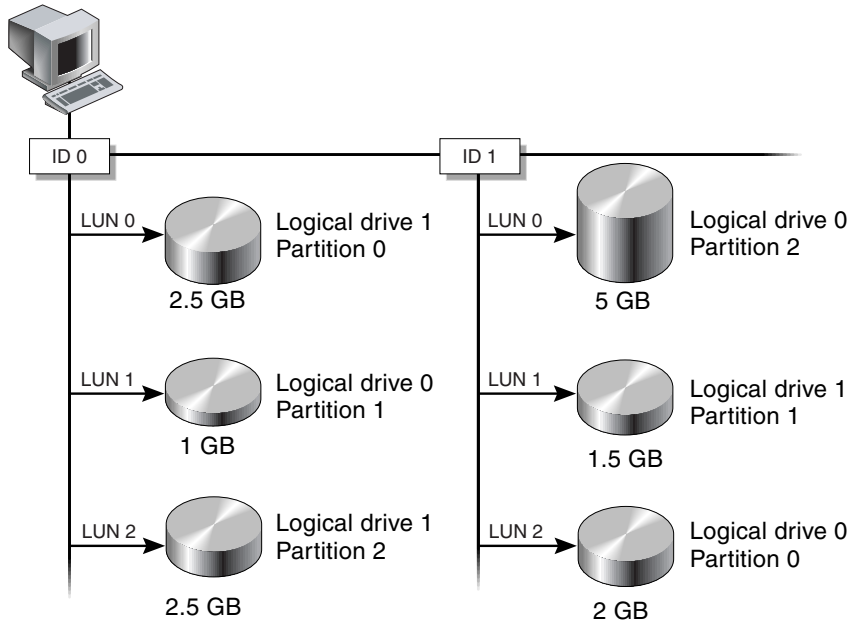


FIGURE A-5 Mapping Partitions to Host ID/LUNs




FIGURE A-6 Mapping Partitions to LUNs Under an ID

RAID Levels

There are several ways to implement a RAID array, using a combination of mirroring, striping, duplexing, and parity technologies. These various techniques are referred to as RAID levels. Each level offers a mix of performance, reliability, and cost. Each level uses a distinct algorithm to implement fault tolerance.

There are several RAID level choices: RAID 0, 1, 3, 5, 1+0, 3+0 (30), and 5+0 (50). RAID levels 1, 3, and 5 are the most commonly used.

The following table provides a brief overview of the RAID levels.

TABLE A-1 RAID Level Overview

RAID Level	Description	Number of Drives Supported	Capacity	Redundancy
0	Striping	2–36	N	No
1	Mirroring	2	N/2	Yes
1+0	Mirroring and striping	4–36 (even number only)	N/2	Yes
3	Striping with dedicated parity	3–31	N-1	Yes
5	Striping with distributed parity	3–31	N-1	Yes
3+0 (30)	Striping of RAID 3 logical drives	2–8 logical drives	N-# of logical drives	Yes
5+0 (50)	Striping of RAID 5 logical drives	2–8 logical drives	N-# of logical drives	Yes

Capacity refers to the total number (N) of physical drives available for data storage. For example, if the capacity is N-1 and the total number of disk drives in the logical drive is six 36-Mbyte drives, the disk space available for storage is equal to five disk drives—(5 x 36 Mbyte or 180 Mbyte). The -1 refers to the amount of striping across six drives, which provides redundancy of data and is equal to the size of one of the disk drives.

For RAID 3+0 (30) and 5+0 (50), *capacity* refers to the total number of physical drives (N) minus one physical drive (#) for each logical drive in the volume. For example, if the total number of disk drives in the logical drive is twenty 36-Mbyte drives and the total number of logical drives is 2, the disk space available for storage is equal to 18 disk drives—18 x 36 Mbyte (648 Mbyte).

RAID 0

RAID 0 implements *block striping*, where data is broken into logical blocks and is striped across several drives. Unlike other RAID levels, there is no facility for redundancy. In the event of a disk failure, data is lost.

In block striping, the total disk capacity is equivalent to the sum of the capacities of all drives in the array. This combination of drives appears to the system as a single logical drive.

RAID 0 provides the highest performance. It is fast because data can be simultaneously transferred to or from every disk in the array. Furthermore, read/writes to separate drives can be processed concurrently.




FIGURE A-7 RAID 0 Configuration

RAID 1

RAID 1 implements *disk mirroring*, where a copy of the same data is recorded onto two drives. By keeping two copies of data on separate disks, data is protected against a disk failure. If, at any time, a disk in the RAID 1 array fails, the remaining good disk (copy) can provide all of the data needed, thus preventing downtime.

In disk mirroring, the total usable capacity is equivalent to the capacity of one drive in the RAID 1 array. Thus, combining two 1-Gbyte drives, for example, creates a single logical drive with a total usable capacity of 1 Gbyte. This combination of drives appears to the system as a single logical drive.

Note – RAID 1 does not allow expansion. RAID levels 3 and 5 permit expansion by adding drives to an existing array.




FIGURE A-8 RAID 1 Configuration

In addition to the data protection that RAID 1 provides, this RAID level also improves performance. In cases where multiple concurrent I/O is occurring, that I/O can be distributed between disk copies, thus reducing total effective data access time.

RAID 1+0

RAID 1+0 combines RAID 0 and RAID 1 to offer *mirroring and disk striping*. Using RAID 1+0 is a time-saving feature that enables you to configure a large number of disks for mirroring in one step. It is not a standard RAID level option that you can select; it does not appear in the list of RAID level options supported by the controller. If four or more disk drives are chosen for a RAID 1 logical drive, RAID 1+0 is performed automatically.

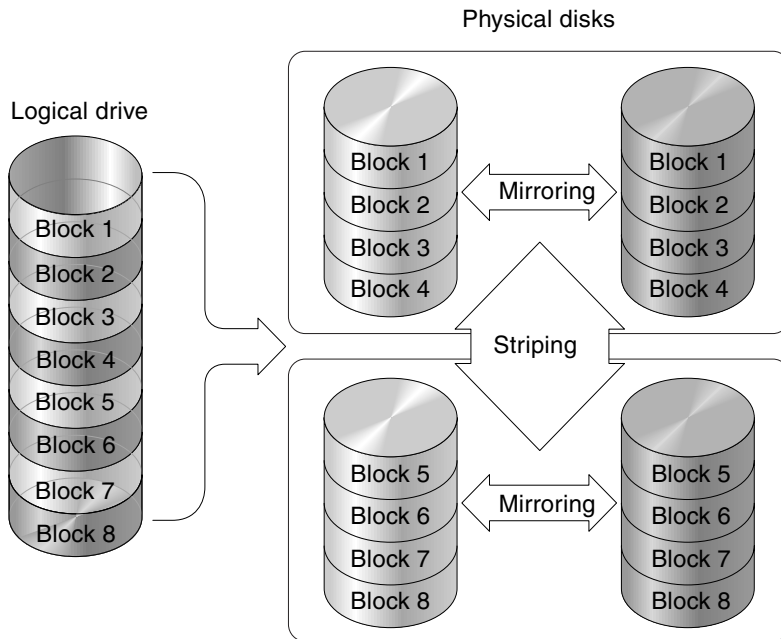


FIGURE A-9 RAID 1+0 Configuration

RAID 3

RAID 3 implements *block striping with dedicated parity*. This RAID level breaks data into logical blocks, the size of a disk block, and then stripes these blocks across several drives. One drive is dedicated to parity. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining disks.

In RAID 3, the total disk capacity is equivalent to the sum of the capacities of all drives in the combination, excluding the parity drive. Thus, combining four 1-Gbyte drives, for example, creates a single logical drive with a total usable capacity of 3 Gbyte. This combination appears to the system as a single logical drive.

RAID 3 provides increased data transfer rates when data is being read in small chunks or sequentially. However, in write operations that do not span every drive, performance is reduced because the information stored in the parity drive needs to be recalculated and rewritten every time new data is written, limiting simultaneous I/O.

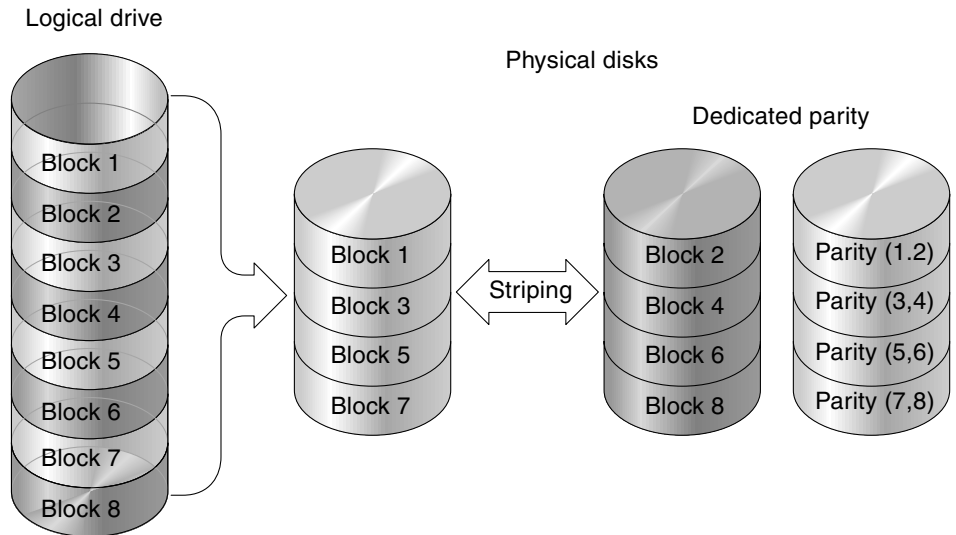


FIGURE A-10 RAID 3 Configuration

RAID 5

RAID 5 implements *multiple-block striping with distributed parity*. This RAID level offers redundancy with the parity information distributed across all disks in the array. Data and its parity are never stored on the same disk. In the event that a disk fails, original data can be reconstructed using the parity information and the information on the remaining disks.




FIGURE A-11 RAID 5 Configuration

RAID 5 offers increased data transfer rates when data is accessed in large chunks, or randomly and reduced data access time during many simultaneous I/O cycles.

Advanced RAID Levels

Advanced RAID levels require the use of the array's built-in volume manager. These combination RAID levels provide the protection benefits of RAID 1, 3, or 5 with the performance of RAID 1. To use advanced RAID, first create two or more RAID 1, 3, or 5 arrays, and then join them. The following table provides a description of the advanced RAID levels.

TABLE A-2 Advanced RAID Levels

RAID Level	Description
RAID 3+0 (30)	RAID 3 logical drives that have been joined together using the array's built-in volume manager.
RAID 5+0 (50)	RAID 5 logical drives that have been joined together using the array's volume manager.

Local and Global Spare Drives

The external RAID controllers provide both local spare drive and global spare drive functions. The local spare drive is used only for one specified drive; the global spare drive can be used for any logical drive on the array.

The local spare drive always has higher priority than the global spare drive. Therefore, if a drive fails and both types of spares are available at the same time or a greater size is needed to replace the failed drive, the local spare is used.

If there is a failed drive in the RAID 5 logical drive, replace the failed drive with a new drive to keep the logical drive working. To identify a failed drive, refer to the *Sun StorEdge 3000 Family RAID Firmware User's Guide* for your array.



Caution – If, when trying to remove a failed drive, you mistakenly remove the wrong drive, you can no longer access the logical drive because you have incorrectly failed another drive.

A local spare drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the local spare drive becomes a member drive and automatically starts to rebuild.

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and there is a local spare and a global spare drive available, the local spare drive is used.




FIGURE A-12 Local (Dedicated) Spare

A global spare drive is available for all logical drives rather than serving only one logical drive (see [FIGURE A-13](#)). When a member drive from any of the logical drives fails, the global spare drive joins that logical drive and automatically starts to rebuild.

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and there is a local spare and a global spare drive available, the local spare drive is used.




FIGURE A-13 Global Spare

Having Both Local and Global Spares

In [FIGURE A-14](#), the member drives in logical drive 0 are 9-Gbyte drives, and the members in logical drives 1 and 2 are all 4-Gbyte drives.




FIGURE A-14 Mixing Local and Global Spares

A local spare drive always has higher priority than a global spare drive; that is, if a drive fails and both a local spare and a global spare drive are available, the local spare drive is used.

In [FIGURE A-14](#), it is not possible for the 4-Gbyte global spare drive to join logical drive 0 because of its insufficient capacity. The 9-Gbyte local spare drive aids logical drive 0 once a drive in this logical drive fails. If the failed drive is in logical drive 1 or 2, the 4-Gbyte global spare drive immediately aids the failed drive.

Monitoring JBODs

This appendix explains how to enable and monitor a standalone JBOD. Note that not all of the array functionality is supported for JBODs. Procedures covered in this appendix include the following:

- [“To Enable JBOD Support” on page 189](#)
- [“To View Component and Alarm Characteristics” on page 191](#)
- [“To Download Firmware for Devices” on page 191](#)
- [“To Discover a Drive” on page 192](#)

Note – Most Sun StorEdge 3000 family products offer a standalone JBOD. The Sun StorEdge 3120 SCSI array is a standalone JBOD. Only the procedures contained in this appendix apply to the Sun StorEdge 3120 SCSI array.

▼ To Enable JBOD Support

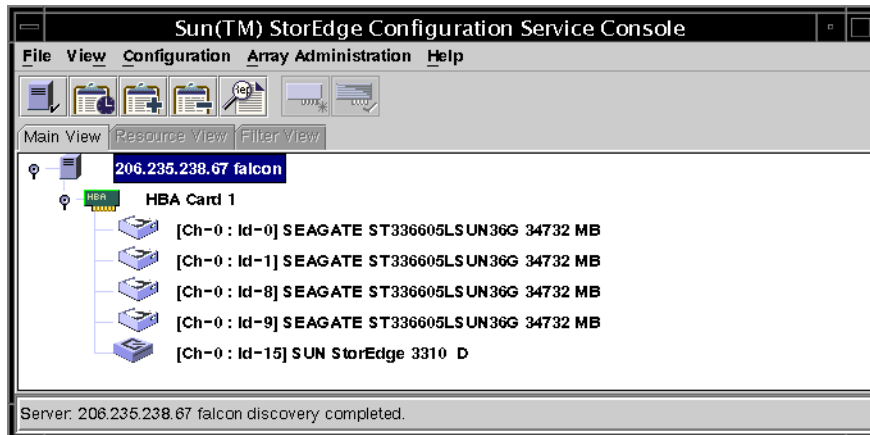
To monitor peripheral device condition and events of a JBOD, you first need to enable JBOD support.

1. **Choose View → Agent Options Management.**
The Agent Options Management window is displayed.
2. **Select the Enable JBOD Support check box.**
3. **To immediately display the JBOD in the main window, you need to probe for new inventory. Choose View → View Server and click Probe.**

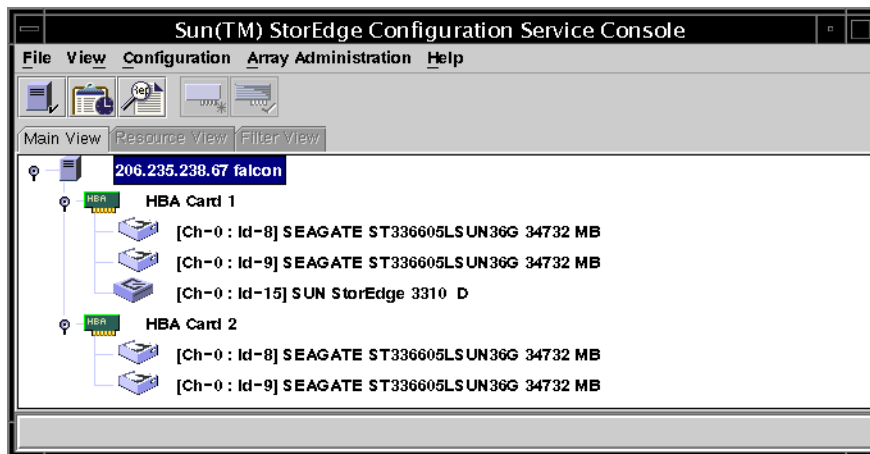
4. Click OK.

The JBOD is displayed in the main window.

In a single-bus configuration, both ports of the JBOD are connected to one HBA on the server as shown in the following example.



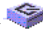
In a dual-bus configuration, each port is connected to its own HBA as shown in the following example.



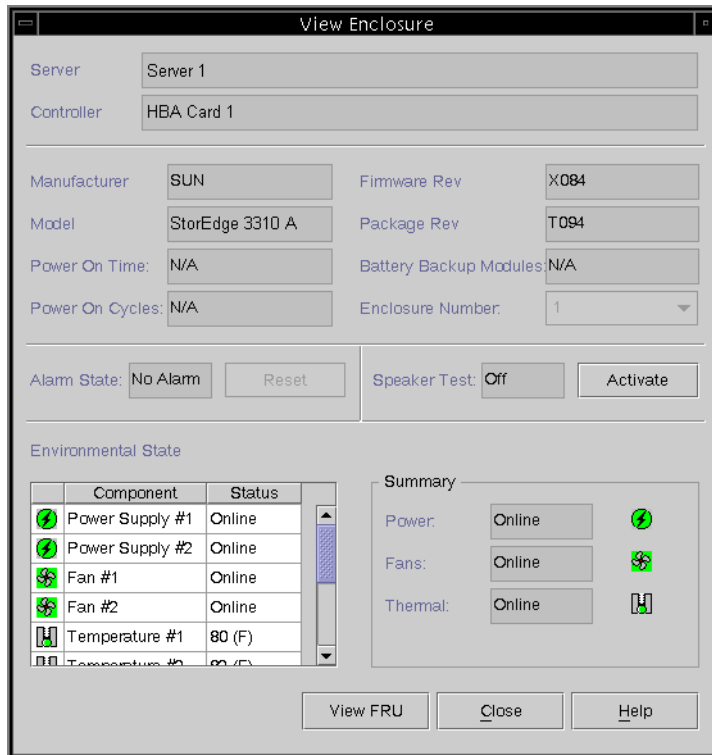
Note – In a dual-bus configuration, if each port is connected to different servers, the program can only monitor the JBOD from the server connected to port B.

▼ To View Component and Alarm Characteristics

The View Enclosure window displays the component and alarm characteristics of a JBOD. When a JBOD is in a state that requires attention, a device status state symbol is displayed on the JBOD device in the main window. See “[Device Status State](#)” on [page 71](#) for a description of the symbols.

1. Select the EMU (SCSI) or SES (Fibre Channel) icon .
2. Choose View → View Enclosure.


To display FRU ID information, click View FRU.



▼ To Download Firmware for Devices

For instructions on how to download firmware to disk drives for the Solaris host, refer to the README file in the patch that contains the firmware. For patch information, refer to the release notes for your array.

▼ To Discover a Drive

Drive failure places the JBOD in a critical state. A red device status symbol  is displayed on the JBOD device in the main window. After replacing a failed drive, perform the following steps to discover the drive.

Solaris OS

The following steps describe how to discover the replaced drive on systems running the Solaris OS.

1. If the drive comes back online automatically (that is, the red device status symbol is dismissed), run the command:

```
# devfsadm
```

2. Run the command:

```
# format
```

3. If the drive does not come back online automatically (that is, the red device status symbol is not dismissed), perform a reconfiguration reboot by running the command:

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

Linux OS

The following steps describe how to discover the replaced drive on systems running the Linux OS.

1. Reboot the system.
2. Run the command:

```
# dmesg
```

3. In the `dmesg` output, look for a line similar to "Detected scsi disk `sdX` at `scsi<controller>`, `id <channel>`, `lun <target>`" where the `X` in `sdX` is the disk number.
4. To create the device entries in `/dev`, run the command:

```
# cd /dev; ./MAKEDEV sdX
```

Where the `X` in `sdX` is the disk number.

5. Proceed to `fdisk`, format, and create a file system.

Microsoft Windows OS

The following steps describe how to discover the replaced drive on a Microsoft Windows OS.

1. Install the new drive according to the instructions in the *Sun StorEdge 3000 Family Installation, Operation, and Service Manual* for your array.
2. Select **View** → **View Server**, and click **Rescan**.
3. To verify that the program recognizes the new drive, select it on the main window.
4. Select **View** → **View Physical Drive**, and verify the information.

HP-UX OS

The following steps describe how to discover the replaced drive on systems running the HP-UX OS.

1. Run the command:

```
# ioscan -fnC disk
```

2. Select View → View Server, and click Rescan.
3. To verify that the program recognizes the new drive, select it on the main window.
4. Select View → View Physical Drive, and verify the information.
5. If the drive is still not seen, the host might need to be rebooted. Run the commands:

```
# sync; sync; sync  
# reboot
```

IBM AIX OS

The following steps describe how to discover the replaced drive on systems running the IBM AIX OS.

Note – You must have superuser privileges to run the commands necessary to replace a failed drive.

1. Create the logical drive and map its LUN to the correct host channel.
2. Run the command:

```
# cfmgr
```


3. Run the command:

```
# lspv
```

Output similar to the following is displayed.

```
hdisk0 000df50dd520b2e rootvg  
hdisk1 000df50d928c3c98 None  
hdisk1 000df50d928c3c98 None
```

4. If any of the drives show “none,” you need to assign a Physical Volume IDENTIFIER.

5. Run the command:

```
# smitty
```

- a. Select Devices.
 - b. Select Fixed Disk.
 - c. Select Change/Show Characteristics of a Disk.
 - d. Select the disk without a pvid.
 - e. Select ASSIGN physical volume identifier, press Tab once to display Yes for the value, and press Return.
 - f. Press Return again to confirm and repeat steps a-g as necessary.
6. From the smitty main menu, select System Storage Management (Physical & Logical Storage) → Logical Volume Manager → Volume Groups → Add a Volume Group.
7. Specify a name for the volume group, make sure the partitions for the journaled file system are large enough, and select the Physical Volume Name(s).
8. From the smitty main menu, select System Storage Management (Physical & Logical Storage) → File Systems → Add / Change / Show / Delete File Systems → (Enhanced) Journaled File System.
9. Select the volume group and set the field.

Run the command:

```
# umount mount point
```


Using the Cluster Configuration (SCSI Only)

For the Sun StorEdge 3310 SCSI array, if you are using Windows NT, Sun StorEdge Configuration Service can monitor shared SCSI storage. This section covers the following topics:

- [“Planning the Cluster Configuration” on page 197](#)
- [“Cluster Configuration Requirements” on page 198](#)
 - [“To Set Up Cluster Configuration” on page 198](#)

A cluster configuration must be set up with the appropriate hardware first, and it also requires Windows NT Server, Enterprise Edition, with Service Pack 3 or 4, and Microsoft Cluster Server (MSCS) software.

If you are planning to set up a cluster configuration, make sure to read this chapter.

Planning the Cluster Configuration

Before you set up a cluster configuration, you need to determine what type of configuration you want because it can make a difference in how you configure the storage initially.

There are two main types of cluster configuration:

- A hot standby server cluster where all of the LUNs belong to one of the two servers in the cluster. If the server owning the LUNs fails, the LUNs are moved to the second server, which has been idle to this point. A minimum of two LUNs are required in this configuration.
- A load balanced cluster, where some LUNs are maintained on one server and some LUNs are maintained on the other server. Both servers process data at the same time, but they are processing I/O on different LUNs.

A minimum of three LUNs must be defined. This enables one small LUN to be established for quorum disk use and a large LUN to be used on each server in the cluster. The quorum disk maintains the cluster configuration data necessary to recover the cluster in the event of a server failure.

Cluster Configuration Requirements

When you are installing the MSCS software, identify the disk to be used for the quorum disk, which maintains cluster information.

In a cluster configuration, the program runs on only one server at a time, the server that has the quorum disk. If the server running with Sun StorEdge Configuration Service malfunctions, the Cluster Administrator automatically changes the first server's disk load to the second server and start the services on that server.

In a two-server cluster configuration, the cluster itself, with its own IP address, becomes the managed server on the Managed Servers list. Add the cluster to the Managed Servers list as the last step in setting up the cluster.

▼ To Set Up Cluster Configuration

Following is an overview of the steps needed to set up an array in a cluster configuration with two host servers.

1. **Set up the servers:**
 - a. **Set up the two servers, each with a PCI bus and a boot disk that is not a part of the shared SCSI storage.**
 - b. **Install a host adapter card such as Ultra-Wide, differential host adapter, in each server and set a unique SI ID for each host adapter on the shared SCSI bus.**
 - c. **Install two network adapter cards on each server.**
 - d. **Install the Windows NT Server, Enterprise Edition, software and Windows NT Service Pack 3, 4, 5, or 6 on the boot disk of each server.**
2. **Install the dual active-active storage subsystem and connect it to both host adapters.**

Refer to the supplied documentation for installation steps.
3. **Install the agent on each server:**
 - a. **Stop the services on one server before installing it on the other server.**

- b. **Ensure that the service is running on a server that has access to a host LUN mapped on logical drive assigned to primary controller.**
- c. **See the appropriate installation chapter for specific steps for installing the agent.**

After installing the agent, you do not need to reboot. However, once the agent is installed, be sure to stop the services on one of the servers.

Note – In the next few steps, you will work with one server only as an example.

4. Install the console software.

You can install the console on one of the servers or on any computer on the network where the servers reside. The program enables you to configure and monitor the array remotely. See the appropriate installation chapter for detailed steps.

5. Add the server that has started the services to the console’s Managed Servers list (see “To Add Servers” on page 12).

Be sure to select Auto Discovery and add the `ssmon` password.

6. Use the console software to verify and configure the storage on the active server and then reboot that server.

Your storage array might already be preconfigured on the dual active–active storage subsystem. You need to look at the configuration in the tree view to determine whether this is the case.

If the storage is not configured or you want to change the configuration, configure all of the LUNs on one server. Later, after the MSCS software is installed, you can allocate the storage between the servers with Cluster Administrator.

7. Use NT Disk Administrator to create partitions and format the LUNs on the active server:

a. If necessary, reassign drive letters for the disks.

Note that partitions must be formatted with NTFS.

Windows NT sees the LUNs on the dual active–active controllers as being on both servers in the cluster. You can create partitions and logical drives on only one server. Later, after MSCS is installed, you can use the Cluster Administrator to apportion the storage between the two servers.

The drive letters for the shared SCSI storage for the two servers must be the same. If there is an additional CD-ROM drive or external hard disk on one server and not on the other, you might need to reassign the drive letters for the shared storage. After this step, check them on the other server to make sure they are the same.

b. Write down the drive letters assigned to the shared storage.

8. **Access the second server, start Disk Administrator and make sure that the second server has the same drive letters as the first.**

If it does not, reassign the drive letters so that they are consistent on both servers.

9. **Shut down the second server.**
10. **Install the MSCS software on the first server and reboot.**
11. **Start Cluster Administrator and make sure that it can see the cluster.**
12. **Access the second server, install MSCS on that server by joining it to the first, and reboot.**
13. **Reboot the second server and verify in Cluster Administrator that both servers are included in the cluster.**
14. **Adjust the disk groups in the Cluster Administrator.**

Make sure that the quorum disk and the other disks, if any, are together in one disk group under the first server where you originally configured the storage. Refer to the MSCS documentation for details to perform this step.

Note – Once you put the two LUNs in the same disk group, you have an empty disk group that you can delete.

15. **Add the services to the group that has the quorum disk on the first server:**
 - a. **On both servers, stop each of the three services, Configuration Service Startup, Configuration Service Monitor, and Configuration Service Server, and set them to manual.**
 - b. **Use Cluster Administrator to install each service as a resource to the group that has the quorum disk.**

For each service, enter it in the following order and type its name in the format indicated (with the two words run together). After each service is installed, bring the service online to start it on the active server.

Enter the services as generic services. You are asked to indicate the dependencies for each resource. The dependencies in the group are as follows:

- Disk associated with the quorum disk
- Other disk, if any to be added to this group
- Configuration ServiceStartup
- Configuration ServiceMonitor
- Configuration ServiceServer

Configuration ServiceStartup is dependent on the two disks that are already in the group. Configuration ServiceMonitor is dependent on Configuration Service Startup, and Configuration ServiceServer is dependent on Configuration Service Monitor.

16. For the two cluster servers to be displayed as one icon under the cluster IP address, edit CLUSTER.TXT.

The file is located in the same directory as the console files. It is on the system disk of the computer where the console is installed. If Drive C is the system disk, the path is:

C:\Program Files\Sun\sscs

The following text shows the contents of the file:

```
#Edit this file to map the NT cluster IP address to several
#servers constituting the cluster.
#The format of the entry is:
#<Cluster IP Address>=<server1 IP Address>:<server2 IP Address>
#The IP address must be in dot form.
#Example: If the cluster IP address is 151.239.130.70 and the
#IP addresses of individual servers are 151.239.130.71 and
#151.239.130.72, the entry would be:
# 151.239.130.70=151.239.130.71:151.239.130.72
#
#IMPORTANT NOTE:
#Use only the Cluster IP address to configure an NT cluster
#server on the Configuration Service console.
#
#151.239.130.70=151.239.130.71:151.239.130.72
```

17. Edit the last line of the file to insert the IP address of the cluster and then the IP addresses of the two servers that make up the cluster by removing the number symbol (#).

18. Start the console and remove the first server from the Managed Servers list, and add the cluster as a server to the list.

19. Access the Server List Setup function. See ["To Add Servers" on page 12](#).

Be sure to select Auto Discovery and add the `ssmon` password. If you type in the cluster name, the program supplies the IP address if your network has DNS.

At this point, the cluster configuration is properly configured to work with the program. If the server that has the services running on it malfunctions, its disk groups are transferred to the second server and the Cluster Administrator automatically starts the services on that server.

Note – When one server fails, it can take as long as 15 minutes for the cluster icon in the console to turn from purple to gray and perhaps another 10 minutes before it goes back to purple again.

If you want to have a load balanced configuration with some of the storage running on the other server, you need to use the Cluster Administrator to move one or more disk groups to the other server.

Determining Host Worldwide Names (Fibre Channel Only)

This appendix explains how to determine the host worldwide name (WWN), worldwide node name (WWNN), and worldwide port name (WWPN) of a FC array. Steps include:

- [“To Determine the WWN” on page 203](#)
- [“To Determine the WWNN” on page 205](#)
- [“To Determine the WWPN” on page 206](#)

▼ To Determine the WWN

Prior to using the LUN Filter feature, it is helpful to identify which Sun StorEdge 3510 Fibre Channel array is connected to which HBA card, and the WWN assigned to each card.

Solaris OS

The following steps explain how to determine the WWN on a Solaris host.

1. **If you have a new HBA device installed on your computer, reboot the computer.**
2. **Type the following command:**

```
# luxadm probe
```

3. Scroll down the listing to see the Fibre Channel device(s) and the related WWN(s).



```
falcon# luxadm probe
Found Fibre Channel device(s):
[Node_WWN:200000c0ff100010] Device Type:Disk device
  Logical Path:/dev/rdisk/c6t220000C0FF100010d0s2
[Node_WWN:201000c0ff000010] Device Type:Disk device
  Logical Path:/dev/rdisk/c6t221000C0FF000010d0s2
```

Linux and Microsoft Windows OS

The following steps explain how to determine the WWN on a Linux and Microsoft Windows host.

1. Boot a specific host, and note the BIOS version and HBA card models connected to your host.
2. Access the HBA card's BIOS with the appropriate command (alt-q or control-a are commonly used).

If the host has multiple HBA cards, select the card which is connected to the storage.

3. Scan the card to look for devices attached to it (usually with the Scan Fibre Devices or the Fibre Disk Utility).

The node name (or similar label) is the worldwide name.

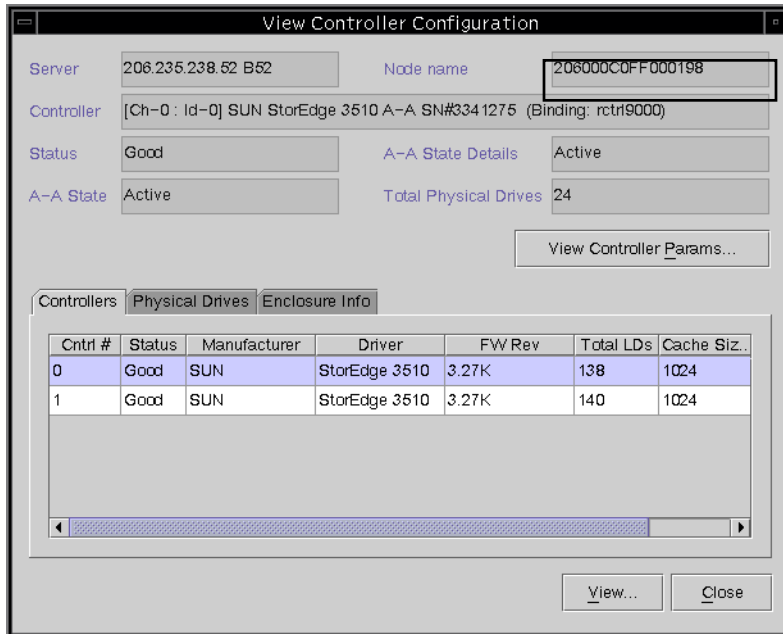
Example with a Qlogic card:

ID	Vendor	Product	Rev	Node Name	Port ID
0	Qlogic	QLA22xx Adapter	B	210000E08B02DE2F	0000EF

▼ To Determine the WWNN


1. Double-click the array icon  in the main Configuration Service window, or select it and choose View → View Controller.

The Node name is displayed in the Node name field in the View Controller Configuration window.

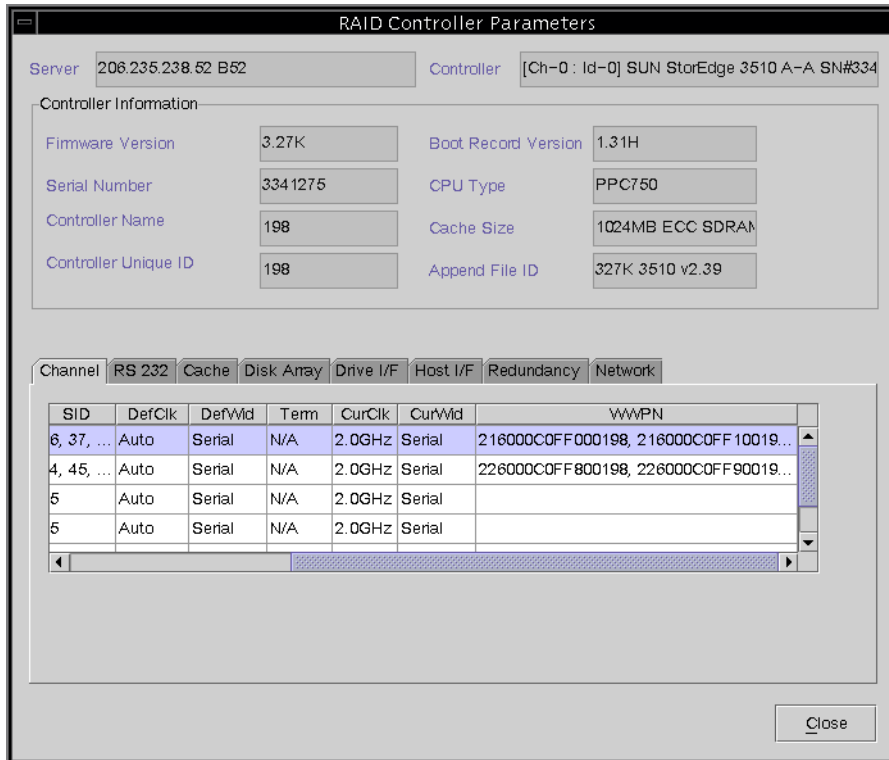


Note – The WWNN information is also included in the xml report generated by the console.

▼ To Determine the WWPN

1. Double-click the array icon  in the main Configuration Service window, or select it and choose View → View Controller.
2. Click View Controller Parameters.

The RAID Controller Parameters window is displayed. On the Channel tab, scroll all the way to the right to display the WWPN.



Note – The WWPN information is also included in the xml report generated by the console.

Email and SNMP

Sun StorEdge Configuration Service has full event monitoring and email notification capabilities. The consoles can send SMTP email messages to a specified list of email addresses. Some email systems, like Microsoft Exchange, can be configured or programmed with scripts to page administration personnel based on events. Some paging services also support sending these email messages to a pager.

The agents place events in the OS error logs. Agents can also send SNMP traps about mass storage events to an enterprise management console, such as HP OpenView. For a list of monitoring software that uses SNMP, visit the Sun Management Center at <http://www.sun.com/software/solaris/sunmanagementcenter/>.

This appendix explains how to set up consoles to send email messages. It also explains how to configure servers to send traps to SNMP management consoles, describes the trap format, and gives background information about SNMP. An additional method for sending SNMP traps that does not require Sun StorEdge Configuration Service is also discussed. Topics covered in this chapter include:

- “How SNMP Works” on page 208
- “Using Sun StorEdge Configuration Service to Send SNMP Traps” on page 212
- “To Send SNMP Traps Without Using Sun StorEdge Configuration Service” on page 221

For details about event monitoring, see “Event Log” on page 94.

How SNMP Works

Simple Network Management Protocol (SNMP) is one of the most widely used protocols for network management. It does what its name says—it manages network devices in a relatively simple manner. SNMP is a simple, unacknowledged, connectionless protocol.

SNMP was originally developed to work under the umbrella of the Internet suite of protocols, primarily in TCP/IP (Transmission Control Protocol/Internet Protocol) transport protocols. Since then, SNMP has been implemented to work over other common transport protocols, such as Novell's IPX/SPX (Internet Packet Exchange/Sequenced Packet Exchange).

SNMP Trap Messages

Agents use SNMP to send information to enterprise management systems.

Once a server is configured to send traps, event information flows to that system.

Trap messages contain the following information:

- OID (Object Identifier) 1.3.6.1.4.1.2294.1.2
- Event Date (MM, DD,YY, such as 01,22,98)
- Event Time (HH:MM:SS, such as 15:07:23)
- Server Address and Name (IP Address Name, such as 192.187.249.187 Administration)
- Card Name
- Event Severity (Informational, Warning, Critical)
- The text message

The information that is displayed and its format are dependent on the SNMP management console you are using.

The file `RST_OID.MIB` must be loaded into the enterprise management console for traps to be received. On a Solaris, Linux, and HP-UX OS, this file is located in `/opt/SUNWsscs/ssagent`. On an IBM AIX OS, this file is located in `/usr/SUNWsscs/ssagent`. On Microsoft Windows, this file is located in `\Program Files\Sun\ssagent`.

Agents and Managers

The SNMP software model consists of agents and managers. An agent is the software that enables variables in the managed device to be monitored. The agent software is either embedded in the device that is managed or running on a computer to which the device is connected.

The manager or management console is the network monitoring software that uses SNMP to send requests to the agents about the variables. The manager polls the agents periodically because the values of variables change. SNMP managers typically analyze the data received from the agents and make it available to the user in a graphical display. The data can be stored in a database for historical and trend analysis.

In addition to responding to the SNMP manager's requests, agents can send unsolicited notifications or traps to the management station. This is similar to interrupts in conventional software.

Management consoles can take various actions in response to a trap. Usually, traps are logged, and the information in them is forwarded through a notification scheme to the interested user. Traps can also cause the SNMP manager to initiate further action, such as querying the agent with more requests to find out the current status of the network entity or launching an application.

Management Information Base (MIB)

On a network, each device instrumented by SNMP agents has one or more variables or objects associated with it. An example of a typical variable might be one that keeps track of the name of the device. These variables or objects are described in a text file called `Management Information Base (MIB)`. A MIB is a precise textual definition of variables being managed. The procedures for defining a MIB have been established by the Internet Engineering Task Force (IETF) of the International Standards Organization (ISO).

SNMP Objects

Some variables for network management are common to all LAN network devices. Many of these variables have been standardized and are referred to as MIB-II variables and MIB-II extensions. The IETF has also defined the procedures for MIBs that are not in a standard MIB-II category. These variables are listed in private enterprise MIBs.

Variables are listed under the entry of the private enterprise MIB. For example, the object ID 1.3.6.1.4.1.2294.1.2 refers to an object ID (or OID) of a variable for keeping track of the agent version under one of Configuration Service's MIBs.

The chart in [FIGURE E-1](#) shows how the OID was derived. The check marks indicate the nodes in the OID. In a MIB, the data of a specific variable or object resides below its OID.

Sun's tree node is 2294. Every private enterprise wanting to develop agents to recognize its own devices can apply to the IETF for its own tree node.

SNMP Request Types

The SNMP operations that support accessing MIB data are *Get*, *GetNext*, and *Set*. *Get* retrieves a scalar value from a table entry field value, given an index value. *GetNext* is used to traverse the instances of table entries, with each command returning the next sequential value in the column (or field) of interest. The *Set* operation is used to create or change the value of a MIB variable.

SNMP Security

Setting values raises the issue of security. SNMP V1 (version 1) has a simple security scheme. Each protocol data unit (PDU) contains a community string, which is like a combination of user name and password. Each agent can be configured for a particular community string. An agent does not respond to a request sent to it except when the community string in the request PDU matches the community string in its own configuration.

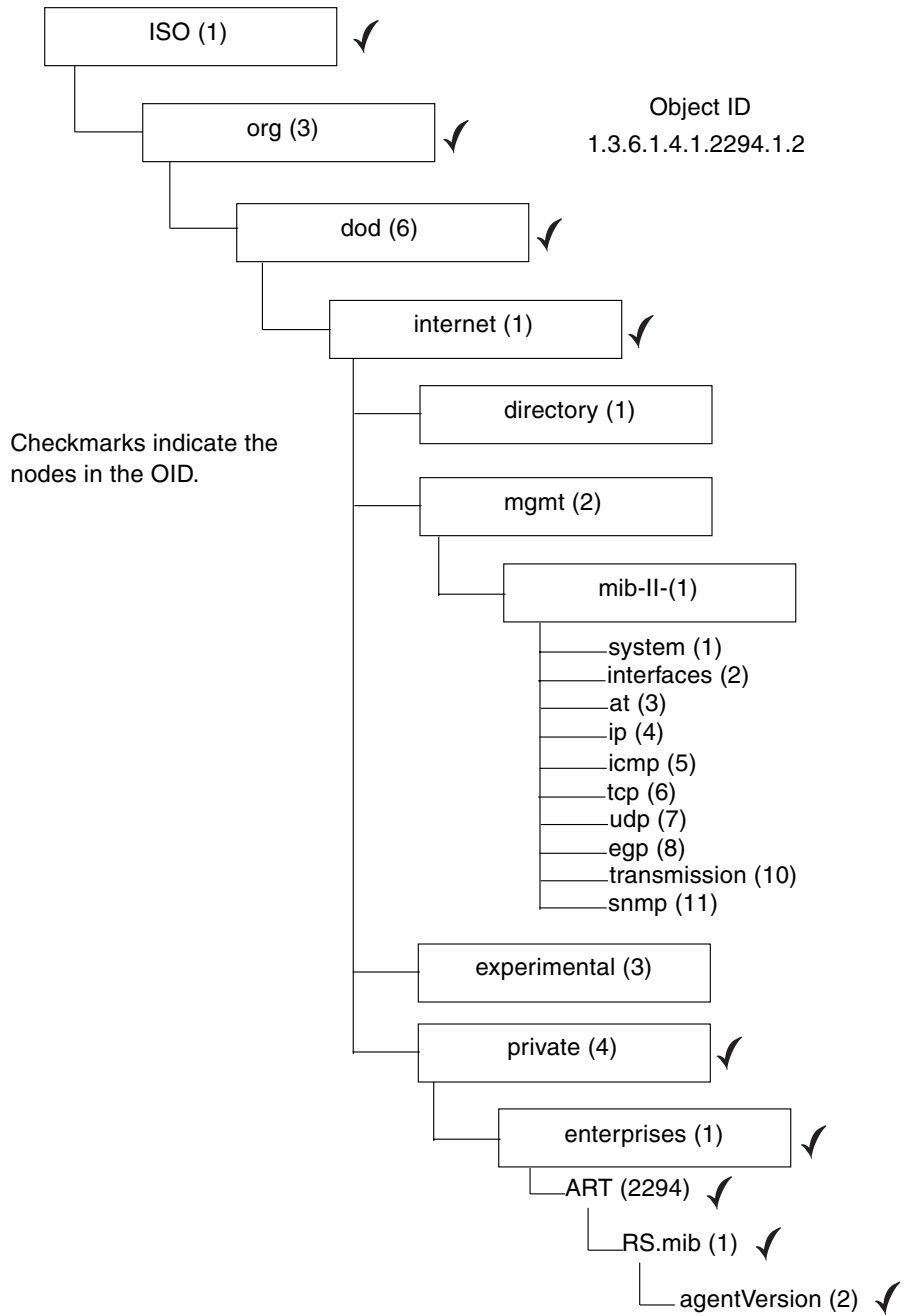


FIGURE E-1 Object Identifier of a MIB Variable

Using Sun StorEdge Configuration Service to Send SNMP Traps


Caution – This section explains how to use Sun StorEdge Configuration Service to send email messages for each server.

▼ To Send Email Messages for Each Server

Configure the console to send email messages for each managed server through the Server List Setup function.

1. Choose File → Server List Setup.

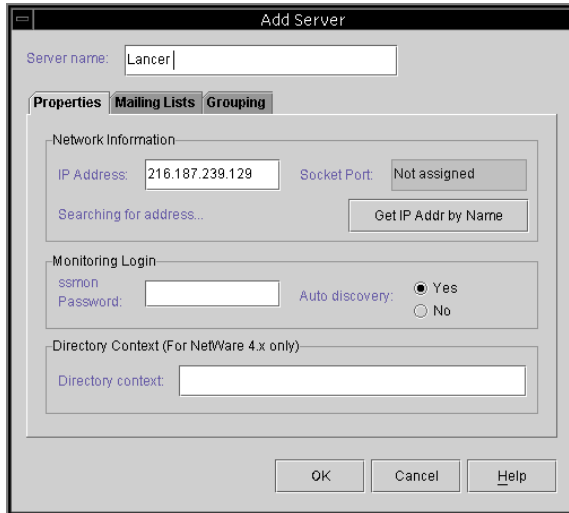
The Server Setup window is displayed.



2. Click Add.

If the server has already been added, you need to move it to the Available Servers list and click Edit.

The Add Server window or Edit Server window is displayed. The two windows have the same fields.

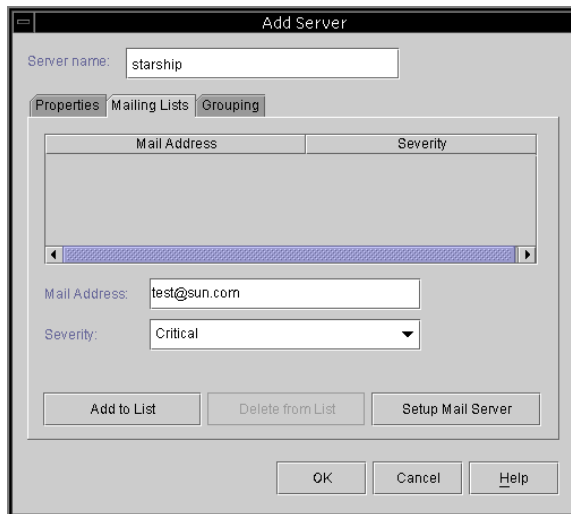


3. If you have not already added the server, fill in the information about the Properties tab.


For detailed information see [“Add a server:”](#) on page 13.

4. To send event messages using email, select the Mailing Lists tab.

The Add Server window is displayed with Mailing Lists selected.



- a. For each user, type an email address in the Mail Address field.
- b. In the Severity list box, scroll through the list and select the desired one.
- c. Click Add to List.
To delete a user from the list, select the mail address and click Delete from List.
- d. Indicate the mail server. Click Setup Mail Server.
The Mail Server Setup window is displayed.



Type the IP address or name of the Simple Mail Transport Protocol (SMTP) mail server that delivers the email messages to the destination addresses specified earlier.

- e. Click OK when you are satisfied with the information about the Mail Server window.
The Add Server or Edit Server window is displayed.
5. Click OK when you are satisfied with the information about the Add Server or Edit Server window.
The Server Setup window is displayed.
6. If you want to add additional servers, repeat Steps 2 through 5 for each server.
7. If necessary, move the server or servers you want this console to control to the Managed Servers list.
If you are editing a server, you must move it back to the Available Servers list.
8. Click OK to return to the main window.

▼ Setting Up Servers to Send Traps

If you have an enterprise management console on your network—for example, HP OpenView or IBM NetView, you can configure servers to send traps about events to the console. This section describes the configuration steps needed.

Note – You need to follow the instructions in this section only if you want to send SNMP traps to other enterprise management consoles.



Caution – If HP OpenView is installed on the same server that you are setting up to send traps, to avoid initialization errors, you need to reconfigure the system resources. For instructions on reconfiguring system resources, see the troubleshooting item [“Symptom: HP OpenView won’t install or the trap daemon won’t start.”](#) on page 227.

Microsoft Windows Servers

To configure a Microsoft Windows server to send SNMP traps to one or more enterprise management consoles, you need to configure each SNMP service on the server to make sure that:

- The server lists the community string or community name of the enterprise management console.
- The server lists the enterprise management console as the recipient of traps from the Microsoft Windows agent software.

▼ To Check the Community String for a Microsoft Windows Host

You can use the default community name or community string of public. However, choosing a community name other than public provides more security because it limits the communication to only those computers that have the specified community string.

1. **For Windows NT 4.0, in the Control Panel window, double-click the Network icon, and then click the Services tab if it is not already selected.**

For Windows 2000, choose Start → Programs → Administrative Tools → Services.

For Windows 2003, choose Start → All Programs → Administrative Tools → Services.

2. **Double-click SNMP Service, and then click Properties.**
3. **The Microsoft SNMP Properties window is displayed.**

For Windows NT, If you do not have the SNMP service installed, you can install it by clicking Add in the Network window. The Select Network Service window is displayed, from which you can select SNMP Service.

Note – You might need to insert the Windows NT 4.0 CD to enable the computer to copy the necessary files for the SNMP service.

For Windows 2000 or Windows 2003, if you do not have the SNMP service installed:

- a. **Choose Start → Control Panel → Add or Remove Programs.**
- b. **Select Add/Remove Windows Components.**
- c. **Select Management and Monitoring Tools, and click Details.**
- d. **Select Simple Network Management Protocol, and click OK.**
- e. **Click Next.**
- f. **Choose Start → Control Panel → Network Connections.**
- g. **Select Local Connection.**
- h. **Select Properties.**

Note – You might need to insert the Windows 2000 or Windows 2003 CD to enable the computer to copy the necessary files for the SNMP service.

4. **Click the Security tab in the Microsoft SNMP Properties window to display the Security settings.**
5. **Make sure that the community name or community string that you want to use for the server is specified in the Accepted community names list.**

If the community name is not already on the list, click Add to display the SNMP Service Configuration window. Type the new name in the Community Name field, specify the rights in the Community Rights field and click Add in that window. The Security tab is displayed with the new name on the Accepted Community Names list.

Note – Community names or community strings are case-sensitive.

6. (Optional) If desired, change the default setting of **Accept SNMP Packets from Any Host to Only Accept SNMP Packets from These Hosts** and add the network address of the enterprise management computer that receives the traps from the agent on the server.

Using this setting provides more security because it lets you specify the computer or computers from which the Windows NT agents accept information.

To add the network address of the computer with Sun StorEdge Configuration Service console software to the lower list box, click **Add** to access the Security Configuration window. Type the IP address in the **Entry** box and click **Add** in that window. The Security tab reappears with the Configuration Service management console computer's address.

7. Click **OK**.

▼ To Specify the Trap Recipient for a Microsoft Windows Host

In this procedure, you are listing the enterprise management console as a recipient of traps from the server.

Note – For the procedure, you need the IP address of the enterprise management console that receives the traps. The procedure also requires the same community name as specified in the preceding sections.

The steps for specifying an enterprise management console computer as the recipient of the agent traps are as follows.

1. Click the **Traps** tab of the **Microsoft SNMP Properties** window to display the **Traps** settings.
2. Make sure the same community name is specified as the community name that was indicated in the **Security** tab earlier.
If you need to correct the community name, type the name in the **Community Name** field and click **Add**. If another name is already in the entry field, it is retained after you type over it. To remove a name, select it and click **Remove**.
3. Add the IP address of the enterprise management console in the **Trap Destinations** list:

- a. Click **Add**.

The Service Configuration window is displayed.

- b. Type the IP address in the **Entry** field and click **Add**.

The Traps tab is displayed with the new network address in the **Trap Destinations** list.

c. Type each IP address in the following format: AAA.BBB.CCC.DDD

You do not need to type leading zeros. An example of an appropriate address is:
192.168.100.1

4. If you want to send optional traps to another enterprise management console, type the network address of that management workstation.

If you have more than one enterprise management console, each one needs to have its address indicated here.

5. Click OK to confirm your settings.

6. Click Close to leave the Network window.

7. Exit Microsoft Windows and restart the server.

If you are planning to install Sun StorEdge Configuration Service, you do not have to reboot until after you install the software.

▼ To Set Up a Solaris Host

For a Solaris host, you can generate SNMP version 1 traps through an interface to an SNMP trap daemon called `sstrapd`. By default, this daemon does not start automatically during the boot process. The following steps describe how to enable trap processing.

1. Create the file `/var/opt/SUNWsscs/ssagent/sstrapd.conf` using any standard text editor.

This file contains the name or the IP address of the SNMP manager console or consoles. If there is more than one console, list them on separate lines.

2. Edit the file `/etc/init.d/ssagent` and remove the comment marks from the SNMP-related start section. These lines are marked by an initial double hash mark (`##`).

3. Save the file after these edits.

The `sstrapd` daemon starts at the next boot or can be started immediately by running the following command

```
/etc/init.d/ssagent uptrap
```

This starts the `sstrapd` daemon as well as the other two daemons if they are not currently running. Only one instance of each daemon is running at this point, whether or not any daemons were previously running.

▼ To Set Up a Linux Host

For a Linux host, you can generate SNMP version 1 traps through an interface to an SNMP trap daemon called `sstrapd`. By default, this daemon does not start automatically during the boot process. The following steps describe how to enable trap processing.

1. **Create the file `/var/opt/SUNWsscs/ssagent/sstrapd.conf` using any standard text editor.**

This file must be created with a list of either system names or IP addresses (one per line) for the SNMP manager consoles. The file can contain blank lines and comment lines.

2. **Edit the file `/etc/init.d/ssagent` and remove the comment marks from the SNMP-related start section.**

These lines are marked by an initial double hash mark (`##`).

3. **Save the file after these edits.**

The `sstrapd` daemon starts at the next boot or can be started immediately by running the following command

```
/etc/init.d/ssagent uptrap
```

This starts the `sstrapd` daemon as well as the other two daemons if they are not currently running. Only one instance of each daemon is running at this point, whether or not any daemons were previously running.

▼ To Set Up an HP-UX Host

1. **Create the file `/var/opt/SUNWsscs/ssagent/sstrapd.conf` using any standard text editor.**

This file must be created with a list of either system names or IP addresses (one per line) for the SNMP manager consoles. The file can contain blank lines and comment lines.

2. Edit the file /sbin/init.d/ssagent using any standard text editor.

Change the following line

```
# Look at environment variable from /etc/rc.config.d/ssagent to if  
["$SSTRAPD"]=1]; thenP trap daemon sstrapd:
```

to

```
# Look at environment variable from /etc/rc.config.d/ssagent to if  
["$SSTRAPD"]=0]; thenP trap daemon sstrapd:
```

3. Save the file after these edits.

The sstrapd daemon starts at the next boot or can be started immediately by running the following command

```
/sbin/init.d/ssagent start
```

This starts the sstrapd daemon as well as the other two daemons if they are not currently running. Only one instance of each daemon is running at this point, whether or not any daemons were previously running.

▼ To Set Up an IBM AIX Host

For an IBM AIX host, you can generate SNMP version 1 traps through an interface to an SNMP trap daemon called sstrapd. By default, this daemon does not start automatically during the boot process. The following steps describe how to enable trap processing.

1. Create the file /var/opt/SUNWsscs/ssagent/sstrapd.conf using any standard text editor.

This file must be created with a list of either system names or IP addresses (one per line) for the SNMP manager consoles. The file can contain blank lines and comment lines.

2. Edit the file /etc/ssagent and remove the comment marks from the SNMP-related start section.

These lines are marked by an initial double hash mark (##).

3. Save the file after these edits.

The `sstrapd` daemon starts at the next boot or can be started immediately by running the following command

```
/etc/ssagent uptrap
```

This starts the `sstrapd` daemon as well as the other two daemons if they are not currently running. Only one instance of each daemon is running at this point, whether or not any daemons were previously running.

To Send SNMP Traps Without Using Sun StorEdge Configuration Service

The following is an alternative method for sending SNMP traps that does not require Sun StorEdge Configuration Service.

The array can send SNMP traps about mass storage events to an enterprise management console, such as HP OpenView. The following configuration does not require that the agent be installed in order to send SNMP traps.

1. Open a text editor.

2. Type the following information, replacing the variables as they apply to your network:

```
[SNMP_TRAP]
ENABLED=1
RECEIVER1=123.123.123.123

[EMAIL]
ENABLED=1
SUBJECT=RAID Event
SENDER_MAIL_BOX=someone@somecompany.com
SMTP_SERVER=123.123.123.123
RECEIVER1=somebody@somecompany.com
```

3. Save the file using the following name `agent.ini`.

4. FTP from your workstation to the array.

5. Type the user name of root and press Return

6. When prompted for a password, press Return.

If you have previously specified a password for the controller, you need to type that password when prompted.

7. Once logged in, change to the `/cfg` directory.

8. Put the file from your workstation into the `/cfg` directory.

9. Quit the FTP session.

10. Reset the controllers through which ever method you are using to manage the array, for example, console, Telnet, or terminal session.

Troubleshooting

This chapter provides troubleshooting suggestions for the following symptoms:

- “Symptom: A server is not responding or a server may be down.” on page 224
- “Symptom: No logical drives are listed when trying to add a logical volume from existing logical drives.” on page 225
- “Symptom: An IP address of a server in a DHCP environment has changed.” on page 225
- “Symptom: Traps are not received from a server.” on page 226
- “Symptom: HP OpenView won’t install or the trap daemon won’t start.” on page 227
- “Symptom: The console does not show changes when hardware is added or replaced.” on page 228
- “Symptom: Logical drive is not present on a Solaris host.” on page 228
- “Symptom: Environmental alarms are not being reported.” on page 228
- “Symptom: Cannot silence the alarm.” on page 229
- “Symptom: The console appears to run slowly.” on page 229
- “Symptom: Sun StorEdge Diagnostic Reporter stops working.” on page 229
- “Symptom: (UNIX OS) The online help is not displayed.” on page 230

Symptom: A server is not responding or a server may be down.

Make sure the services are running at the server.

1. (UNIX OS) Run the following command:

```
# ps -e | grep ss
```

The names `ssmon` and `ssserver` should both be present in the output. If they are not, proceed to Step 2. If they are, proceed to Step 4.

(Windows NT) Choose Settings → Control Panel → Services and verify that SUNWscsd Monitor, SUNWscsd Server, and SUNWscsd Startup services are started. If they are not, proceed to Step 2. If they are, proceed to [Step 4](#).

(Windows 2000) Choose Start → Programs → Administrative Tools → Computer Management. Click Services & Applications → Services and verify that SUNWscsd Monitor, SUNWscsd Server, and SUNWscsd Startup services are started. If they are not, proceed to Step 2. If they are, proceed to Step 4.

(Windows 2003) Choose Start → Administrative Tools → Computer Management → Services & Applications → Services and verify that SUNWscsd Monitor, SUNWscsd Server, and SUNWscsd Startup services are started. If they are not, proceed to Step 2. If they are, proceed to Step 4.

2. (UNIX OS) Stop and start the daemons as explained in the *Sun StorEdge 3000 Family Software Installation Guide*.

(Microsoft Windows) Stop and start the services by right-clicking the service you want to start or stop.

3. If the daemons/services do not stop/start properly, do a reconfiguration reboot.

4. Make sure that the TCP/IP protocol software is loaded, and that the network card and the network connection are working.

To test the TCP/IP connectivity at the command-line, type:

```
# ping {IP address of the server or the server name}
```

If you do not get a reply, there is a problem in the TCP/IP protocol services. Contact your MIS department for assistance.

5. Make sure that the correct server name and password are specified.

If the name or password is not correct, make the correction by editing the entries. See the “Administrative (User) Security Levels and Guidelines” section for your OS in the *Sun StorEdge 3000 Family Software Installation Guide*.

The password has to be the one established through the procedure for setting up users.

- Make sure the users (*ssmon*, *ssadmin*, *ssconfig*) were set up properly. Be sure that Users are set up in Windows NT with Permissions Act as a part of the operating system and Log on as a batch job.
- If you are using a naming service (Windows NT domains, and NIS or NIS+ in a UNIX OS), make sure that the users have been correctly added to the naming service.
- If you have entered the user names in a domain under Windows NT, allow some time for the user names to propagate.

Symptom: No logical drives are listed when trying to add a logical volume from existing logical drives.

If you do not see any logical drives listed under Select a Logical Drive, the logical drives have not been unmapped and therefore are unavailable to select. You must unmap the logical drives first.

Symptom: An IP address of a server in a DHCP environment has changed.

In a DHCP environment, there is a remote possibility of a server’s IP address changing if the server has been offline for longer than three days.

In these cases, the console must be notified of the new IP address of the server because the console communicates with the server through its IP address.

To remedy this situation, the server’s new IP address needs to be entered into the Edit Server window of the console software.

1. **Choose File → Server List Setup.**
2. **Move the server name from the Managed Servers list to the Available Servers list.**
The Edit Server window is displayed. This window shows the server name and the IP address as it was configured.
3. **Select the name of the server in the Available Servers list, and click Edit.**
The Edit Server window is displayed. This window shows the server name and the IP address as it was configured.

4. If the network address has changed, click Get IP Addr by Name.

The program searches for and displays the correct IP address if you provided the name of the server as it is recorded by the service used by your network. Otherwise, you must type the IP address.

If the name used for the server is not the same as the server's network name, or if the naming service is not yet updated, you have to type the IP address manually.

5. Once the server IP address is entered, click OK to specify the changed address and go back to the Server List Setup window.

6. Move the server name back to the Managed Servers list.

7. Click OK to leave the Server List Setup window.

Symptom: Traps are not received from a server.

To troubleshoot why a trap from a server is not received on an SNMP management workstation, such as HP OpenView, use the following steps:

1. Make sure that the SNMP service has been started on the server.

Verify that `sstrapd` is running by running the following command

```
# ps -e | grep ss
```

The output includes the name `sstrapd`; if it does not, start or stop the agent as explained in the *Sun StorEdge 3000 Family Software Installation Guide*.

Under Windows NT, choose Control Panel → Services and check whether the SNMP service is listed and is shown as started. If it is not, install it, and then start the service. Refer to your Windows NT documentation for details.

2. Check that the proper target address of the enterprise management console and the community string is set up on the server properly.

The detailed steps for the proper setup on a Windows NT server for sending traps to an enterprise management console are given in [“Setting Up Servers to Send Traps” on page 215](#).

3. Make sure that the MIB is compiled under the SNMP management console.

On a Solaris, Linux, and HP-UX OS, the `RST_OID.MIB` file is located in `/opt/SUNWsscs/ssagent`. On an IBM AIX OS, this file is located in `/usr/SUNWsscs/ssagent`. On Microsoft Windows, this file is located in `\Program Files\Sun\ssagent`. Refer to your SNMP management console documentation (such as HP OpenView) for information about how to compile the MIB.

4. Under Windows NT, check the NT registry entries for Sun StorEdge Configuration Service.

If the entries are absent, reinstall the agent on the server.

a. First, make sure that

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SNMP\
Parameters\ExtensionAgents

contains an entry for SSAGENT.

b. Second, check that:

HKEY_LOCAL_MACHINE\SOFTWARE\Sun\Sun StorEdge Configuration
Service Agent\SNMP Extension Agent\Current Version

points to the extended agent program with the complete path name:

C:\PROGRA~1\Sun\ssagent\SSAGENT.DLL

C:\PROGRA~1\Sun\ssagent\
is the default installation directory.

is the default installation directory.

If you have installed the program in another directory, make sure to point it to that directory.

c. Third, make sure that:

SSAGENT.DLL

is in the C:\PROGRA~1\Sun\ssagent directory.

Symptom: HP OpenView won't install or the trap daemon won't start.

If you installed HP OpenView on the same server that you set up to send SNMP traps, there are insufficient system resources for both the trap daemon and HP OpenView to run. Reconfigure system resources as explained in the following steps.

1. Add the following lines to the end of /etc/system:

```
set shmsys:shminfo_shmmax=0x2000000
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=256
set shmsys:shminfo_shmseg=256

set semsys:seminfo_semmap=256
set semsys:seminfo_semmni=512
set semsys:seminfo_semmns=512
set semsys:seminfo_semmsl=32
```

Note – If, due to the requirements of another application installed on your system, `/etc/system` already contains statements assigning values to any of these kernel parameters, you need to merge the parameter assignments shown in Step 1 into the file so that each parameter is only assigned once. If a previous parameter value is different than one specified in Step 1, specify the larger value.

2. Reboot the server.

Symptom: The console does not show changes when hardware is added or replaced.

If you add new equipment or if you replace a failed device, such as a disk drive, tape drive, power supply, or fan, the console does not always show updated information. At times, the tree view and the other views associated with it do not reflect the actual status of a device.

If a new device is added and you want this device to be recognized before the next periodic scan, click Rescan in the Server View window. The server immediately rescans for the inventory, and the console updates its server inventory and present the result of any changes in the tree view. To ensure a timely inventory, do a manual scan.

You might have to click Rescan more than once. When the server is in the process of a rescan and the console sends a request for the inventory, the server sends the console the server inventory at its last update, because it has not yet finished the current scan.

Symptom: Logical drive is not present on a Solaris host.

If the logical drive is not being displayed, be sure the logical drive is labeled and not excluded by Sun StorEdge Configuration Service (if installed).

Symptom: Environmental alarms are not being reported.

Storage enclosures with SCSI-based enclosure monitoring capabilities (such as SAF-TE cards) send environmental alarms to the console. The alarm state can be caused by a failure of a drive, fan, power supply, or battery, or by an abnormal temperature.

If an environmental alarm occurs, you must click Reset in the View Enclosure window to clear the alarm. You need to have configuration security privileges to click Reset in the software.

Symptom: Cannot silence the alarm.

Storage enclosures with SCSI-based enclosure monitoring capabilities (such as SAF-TE cards) send environmental alarms to the console. The alarm state can be caused by a failure of a drive, fan, power supply, or battery, or by an abnormal temperature. To silence an environmental alarm, you must push the Reset button on the right ear of the array.

An alarm can also be caused by a controller event, such as when a logical drive fails, during a rebuild, or when adding a SCSI drive. Refer to the “Event Messages” appendix in the *Sun StorEdge 3000 Family RAID Firmware User’s Guide* for more information about controller events. To silence the controller, see [“To Mute the Controller Beeper” on page 167](#).

Note – Pushing the Reset button has no effect on controller event alarms and muting the beeper has no effect on failed component alarms.

Symptom: The console appears to run slowly.

Sun StorEdge Configuration Service can monitor and manage up to 32 arrays at one time. However, console response time can decrease as the number of arrays being managed increases.

The console under Windows NT requires at least 20 Mbyte of free virtual memory (physical memory plus the page file size) on the console workstation. The console program is more responsive with larger available physical memory, which minimizes the OS *paging* of memory from physical memory to the page file.

When the amount of memory used approaches the total available virtual memory (physical memory plus the page file size), the excessive paging likely causes problems, resulting in poor performance of all applications on that workstation.

Increase the physical memory and page file size to increase the overall virtual memory. To change the page file size, select Control Panel → System, and then choose the Performance tab in the System Properties window that is displayed. In Windows NT, change the page file size under the Virtual Memory section of the Performance tab window.

Symptom: Sun StorEdge Diagnostic Reporter stops working.

(UNIX OS) There are three conditions under which Sun StorEdge Diagnostic Reporter stops working and does not report its condition. The workaround is to stop and restart it.

- If the agent fails or is stopped and restarted, Sun StorEdge Diagnostic Reporter stops working.
- If the Config Tool is running and the daemon is stopped and restarted, a condition might occur whereby the Config Tool can no longer communicate with the daemon.
- If the agent fails or is stopped, the daemon does not detect it, stops sending email messages, and continues to show that Sun StorEdge Diagnostic Reporter is still connected by displaying a green status.

For a Solaris host and Linux host, stop and restart Sun StorEdge Diagnostic Reporter by typing:

```
# /etc/init.d/ssdgrptd stop
# /etc/init.d ssdgrptd start
```

For an HP-UX host, stop and restart Sun StorEdge Diagnostic Reporter by typing:

```
# /sbin/init.d/ssdgrptd stop
# /sbin/init.d ssdgrptd start
```

For an IBM AIX host, stop and restart Sun StorEdge Diagnostic Reporter by typing:

```
# /usr/sbin/ssdgrptd start
# /usr/sbin/ssdgrptd stop
```

Symptom: (UNIX OS) The online help is not displayed.

Make sure the absolute path name of the web browser you want to use to display online help has been specified.

1. **For the Solaris, Linux, and HP-UX OS, change to** /opt/SUNWsscs/sscsconsole.
For the IBM AIX OS, change to /usr/SUNWsscs/sscsconsole.
2. **Type:**

```
./config_sscon
```

3. **Enter the absolute path name of the web browser.**

Error Codes and Messages

This appendix provides a list of error codes and error and status messages.

Error Codes

Shown and described in the Event Log, an error code is an eight-character code that consists of four fields of two characters each. Some codes might be followed by a dash and another eight-character code, which is for internal use only.

The following tables describe the error codes used in each two-character field. The format of the four fields of two characters is:

SSMMmmEE

where

SS = the Severity Field

MM = the Major Field

mm = the Minor Field

EE = the Error Field

TABLE G-1 Severity Field

Severity Field	Description
01	Critical
02	Warning
03	Informational

TABLE G-2 Major Field

Major Field	Description
01	Not Used
03	Legacy
04	RAIDCard Card message
05	JBOD (Just a bunch of disks) message
06	SSMon monitor message
07	dll/shared library message
08	SSserver message
09	HBA (Host Bus Adapter) message
0A	Server Communications message
0B	Controller message

The Minor field is defined in following table. The definition of the Error Field is dependent upon the definition of the Minor Field in which case the appropriate Table is referenced.

TABLE G-3 Minor Field

Minor Field	Description
01	Unknown
02	See “Error Field: System Drive State” on page 234.
03	See “Error Field: Disk State” on page 234.
04	See “Error Field: SAF-TE State” on page 235.
05	See “Error Field: Tape State” on page 235.
06	See “Error Field: Redundancy State” on page 236.
07	See “Error Field: Internal State” on page 237.
08	See “Error Field: Device State” on page 237.
09	See “Error Field: Initialization State” on page 238.
0A	See “Error Field: Invalid Client Parameter” on page 238.
0B	See “Error Field: Open Transport” on page 239.
0C	See “Error Field: Close Transport” on page 239.
0D	See “Error Field: Memory Allocation” on page 239.
0E	RaidCard fault detected.
0F	See “Error Field: Transport” on page 239.
10	See “Error Field: Command” on page 240.
11	See “Error Field: Command” on page 240.
12	See “Error Field: Main Communications” on page 240.
13	See “Error Field: Communications Async” on page 241.
14	See “Error Field: Communications Security” on page 241.
15	See “Error Field: Timeout” on page 241.
16	See “Error Field: Administration” on page 241.
17	See “Error Field: Firmware” on page 242.
18	See “Error Field: System Shutdown” on page 242.
19	Dynamic Growth fault detected.
1C	See “Error Field: Set Config” on page 243.
1D	See “Error Field: Controller Event” on page 243.

TABLE G-3 Minor Field (*Continued*)

Minor Field	Description
1E	See “Error Field: Drive Side Event” on page 244.
1F	See “Error Field: Host Side Event” on page 245.
20	See “Error Field: Logical Drive Event” on page 245.
23	See “Server Manage/Monitor Event Error” on page 245.

Error Field: System Drive State

TABLE G-4 System Drive State Errors

Error Field	
01	Offline
02	Degraded
03	Online
04	Unknown
05	Parity Check
06	Reconstruction
07	Reconfiguration
08	Dynamic Growth
09	Nonexistent
0A	Initialization

Error Field: Disk State

TABLE G-5 Disk State Errors (*1 of 2*)

Error Field	Description
01	Offline
02	Degraded
03	Online

TABLE G-5 Disk State Errors (2 of 2)

Error Field	Description
04	Unknown
05	SMART
06	Dead

*Error Field: SAF-TE State***TABLE G-6** SAF-TE State Errors

Error Field	Description
01	Power
02	Fan
03	Thermal
04	Alarm
05	Locked
06	Slot
07	Unknown
08	Unable to get SAF-TE information
09	Battery
0A	Invalid Number of Slots
0B	Environmental Data Not Available

*Error Field: Tape State***TABLE G-7** Tape State Errors

Error Field	Description
01	Read
02	Write
03	Hard Error
04	Media Error
05	Read Failure

TABLE G-7 Tape State Errors (*Continued*)

Error Field	Description
06	Write Failure
07	Media Life
08	Not Upgradable
09	Write Protect
0A	Nonremoval
0B	Cleaning Media
0C	Unsupported Format
0D	Snapped Tape
14	Clean Now
15	Clean Periodic
16	Expired Cleaning Media
1E	Hardware A
1F	Hardware B
20	Interface
21	Eject Media
22	Down Load Fail
28	Loader Hardware A
29	Loader Tray Tape
2A	Loader Hardware B
2B	Loader Door
C8	Query Log Failed
C9	Inquire Tape Failed

Error Field: Redundancy State

TABLE G-8 Redundancy State Errors

Error Field	Description
01	Mode OK
02	No Memory for Negotiation
03	Secondary Failure

TABLE G-8 Redundancy State Errors (*Continued*)

Error Field	Description
04	Secondary Removed
05	Detected Removal of Secondary
06	Detected Insertion of Secondary
07	Primary Failure
08	Primary Removed
09	Detected Removal of Primary
0A	Detected Insertion of Primary
0B	Unknown State
0C	Partner Failed
0D	Partner Inserted
0E	Partner Removed

Error Field: Internal State

TABLE G-9 Internal State Errors

Error Field	Description
01	No Memory
02	Semaphore
03	Thread
04	No Devices

Error Field: Device State

TABLE G-10 Device State Errors

Error Field	Description
01	Inquiry
02	Unknown

Error Field: Initialization State

TABLE G-11 Initialization State Errors

Error Field	Description
01	dll Initializing
02	dll Initializing Failed

Error Field: Invalid Client Parameter

TABLE G-12 Client Parameter Errors

Error Field	Description
01	Set Configuration – Configuration Miscompare
02	Set Configuration – Configuration Device Invalid
03	Set Configuration – Enquire Miscompare
04	Set Configuration – Enquire2 Miscompare
05	Incorrect Application Length
06	Command not Supported
07	Invalid Command
08	Set Configuration – General Miscompare
09	Invalid Length
0A	Invalid Card Identifier
0B	Invalid Card Name
0C	Invalid Parameter
0D	Invalid Command for Card Type
0E	Set Configuration – Invalid Additional Parameter
0F	Set Configuration – Block Overlap
10	Set Configuration – Device Information Invalid

Error Field: Open Transport

TABLE G-13 Open Transport Errors

Error Field	Description
01	Open Transport

Error Field: Close Transport

TABLE G-14 Close Transport Errors

Error Field	Description
01	Close Transport

Error Field: Memory Allocation

TABLE G-15 Memory Allocation Errors

Error Field	Description
01	Insufficient Memory
02	Insufficient Memory for Administration Operation

Error Field: Transport

TABLE G-16 Transport Field Errors

Error Field	Description
01	Lock Fault
02	Insufficient Memory
03	Acquire Lock Fault
04	Release Lock Fault
05	Invalid Command
06	Invalid Length
07	Invalid Card Name
08	Invalid Card Identification

TABLE G-16 Transport Field Errors *(Continued)*

Error Field	Description
09	No Cards Found
0A	No Devices Found
0B	Open Fault
0C	Card Name not Found

Error Field: Command

TABLE G-17 Command Code Errors

Error Field	Description
01	Lock Fault

Error Field: Main Communications

TABLE G-18 Main Communications Errors

Error Field	Description
01	Socket Fault
02	Report Fault
03	Thread Fault
04	Lock Fault
05	System Fault

Error Field: Communications Link

TABLE G-19 Communications Link

Error Field	Description
01	Socket Fault

Error Field: Communications Async

TABLE G-20 Communications Async

Error Field	Description
01	Socket Fault
02	Thread Fault
03	Cold Link Fault
04	Send Event Fault

Error Field: Communications Security

TABLE G-21 Communications Security

Error Field	Description
01	Security Violation

Error Field: Timeout

TABLE G-22 Timeout Errors

Error Field	Description
01	Configuration Update
02	Lock Timeout

Error Field: Administration

TABLE G-23 Administration Errors

Error Field	Description
01	Set Customer Name Fault
02	Set Configuration Ended
03	Initialize
04	Initialize Ended
05	Rebuild

TABLE G-23 Administration Errors (*Continued*)

Error Field	Description
06	Rebuild Ended
07	Parity Check
08	Parity Check Ended
09	Set SAF-TE Slot State
0A	Set SAF-TE Perform Slot
0B	Set SAF-TE Send Global
0E	Schedule Parity Check
0F	Schedule Parity Check Ended
10	Set Controller Parameters
11	Firmware Download
12	Consistency Check or Rebuild Ended

Error Field: Firmware

TABLE G-24 Firmware Download Errors

Error Field	Description
01	Not a SAF-TE Device
02	Invalid Data Length
03	Download Failed
04	Check Sum Failed

Error Field: System Shutdown

TABLE G-25 System Shutdown Errors

Error Field	Description
01	System Shutdown Failed

Error Field: RDP Path Change

TABLE G-26 RDP Path Change Errors

Error Field	Description
01	RDP Path Failed
02	RDP Path Restored

Error Field: Set Config

TABLE G-27 Set Config Errors

Error Field	Description
01	Set Configuration Thread creation Failed
02	Get Logical Drive List Command Failed
03	Create Config Command Failed
04	Init Completion Status Command Failed
05	Get Configuration Command Failed
06	Change Volume Config Command Failed
07	Delete Logical Drive Command Failed

Error Field: Controller Event

TABLE G-28 Controller Event Errors

Error Field	Description
01	Controller Reset
02	Controller DRAM Parity Error
03	Redundant Controller Failed
04	Controller Power Supply Failed
05	Controller Fan Failed
06	Controller Temperature Alert
07	Controller UPS AC Power Loss

TABLE G-28 Controller Event Errors (*Continued*)

Error Field	Description
08	Controller Initialization Complete
09	Controller Power Supply Back Online
0A	Controller Fan Back Online
0B	Controller Temperature Normal
0C	Controller UPS AC Power Back On

Error Field: Drive Side Event

TABLE G-29 Drive Side Event Errors

Error Field	Description
01	Drive Channel Select Timeout
02	Drive Channel SCSI Bus Error
03	Drive Channel Unexpected Disconnect
04	Drive Channel Negotiation Error
05	Drive Channel Target Timed Out
06	Drive Channel Parity Error
07	Drive Channel Data Under Or Overrun
08	Drive Channel Undefined Error
09	Drive Channel SCSI Bus Reset Issued
0A	Drive Channel Not Ready Error
0B	Drive Channel Target HW Error
0C	Drive Channel Target Media Error
0D	Drive Channel Unexpected Unit Attention
0E	Drive Channel Unexpected Sense Data
0F	Drive Channel Block Reassign Fail
10	Drive Channel Block Reassign Success
11	Drive Channel SCSI Aborted Command
12	Drive Channel SCSI Channel Fail
13	Drive Channel Smart Fail
14	Drive Channel Scan SCSI Drive

Error Field: Host Side Event

TABLE G-30 Host Side Event Errors

Error Field	Description
01	Host Channel SCSI Bus Reset
02	Host Channel SCSI Bus Device Reset
03	Host Channel Abort Tag Message
04	Host Channel Parity Error
05	Host Channel Reselect Timeout

Error Field: Logical Drive Event

TABLE G-31 Logical Drive Event Errors

Error Field	Description
01	Logical Drive SCSI Drive Fail
02	Logical Drive Initialization Fail
03	Logical Drive Rebuild Fail
04	Logical Drive Initialization Commenced
05	Logical Drive Initialization Completed
06	Logical Drive Rebuild Commenced
07	Logical Drive Rebuild Completed

Server Manage/Monitor Event Error

TABLE G-32 Server Manage/Monitor Event Error

Error Field	Description
01	Server to Manage/Monitor has changed.

Error and Status Messages

Most error and status messages are self-explanatory. [TABLE G-33](#) clarifies terms which are used in messages. [TABLE G-34](#) lists the Error/Status messages.

TABLE G-33 Substituted Values

Value	Description
Channel	Card channel number indicating a SCSI channel. A value of 1 indicates the first channel.
Fan	Fan number.
Enclosure	Enclosure number.
LogicalArray	Logical array number.
LogicalDrive	Logical Drive number.
Lun	LUN number.
Name	A text name.
Power	Power supply number.
Slot	Slot number.
State	State of a logical array or a device or an enclosure in text. The values are: Critical, Online, Offline, Critical Rebuild, Non Existent, Low Battery, Normal.
Target	A target or SCSI ID number.
Temperature	The temperature in centigrade.

TABLE G-34 Error/Status Messages

Error and Status Messages
A SCSI Drive Failed (Ch:%d,Id:%d). Replace the defective drive.
A tape has been left in the autoloader by a previous hardware fault. Insert an empty magazine to clear the fault. If the fault does not clear, turn the autoloader off and then on again. If the problem persists, call the tape drive supplier helpline.
A user named sconfig must be created for configuring SUNWsscs Diag Reporter.
Abort Clone SCSI Drive.
Abort Initialization.
Abort Parity Check.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Abort Rebuild.

Abort/Clear-Queue/Abort-Tag Message received (SCSI Channel %d, SCSI ID %d, Target Lun %d). Possible failure of I/O module, cable, controller, or HBA.

Access denied for one or more servers. Log in as ssconfig user, and then try the operation again.

Access denied for one or more servers. Please log in as an ssadmin user and then try the operation again.

Active trap events is empty.

Add Physical drive %d:%d has started on logical drive *LogicalDrive*.

Add Physical drive has completed on logical drive *LogicalDrive*.

Add SCSI Drive into Logical Drive information.

Agent Name is empty.

An error occurred while getting data from the server.

An HBA card with this WWN already exists.

Another instance of this program is already running!

Array Admin in Progress.

At least one email address needs to be configured.

Attempted to load a cartridge type, not supported by this drive. The cartridge has been automatically ejected. Attempted to load an unsupported tape format.

Attempted to write to a write-protected cartridge. Remove the write-protection or use another tape.

Audible alarm has an unknown status

Audible alarm has been triggered.

Audible alarm has been turned off.

Background check or rebuild operation in progress.

Background initialization in progress.

Background parity check in progress.

Background rebuild operation in progress.

Bad Block Encountered (Ch:%d,Id:%d). Sector number is %u. Informational message.

Battery Expiration Monitoring of (Ch:%d,Id:%d) is not available.

Block Reassignment Failure (Ch:%d,Id:%d). Defective drive. Replace.

Block Reassignment Succeeded (Ch:%d,Id:%d.) Informational message. Drive may be failing.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Cannot accept the incoming connection.

Cannot allocate memory.

Cannot bind TCP port with the opened TCP/IP communication.

Cannot connect with the incoming connection.

Cannot create a communication session for the incoming connection.

Cannot create host LUN filter; no host IDs were selected.

Cannot create standard host LUN mapping, no host IDs were selected.

Cannot find the logical drive that was just created. It may have been aborted.

Cannot get configuration from controller. The operation is canceled.

Cannot get configuration in set configuration.

Cannot listen on the opened TCP/IP communication.

Cannot lock a resource.

Cannot open TCP/IP communication.

Cannot perform IOCTL on the TCP/IP communication.

Cannot query incoming data status of the connection.

Cannot receive data from the connected client.

Cannot report status.

Cannot send data to the connected client.

Change logical drive failure.

Change logical volume failure.

Change or set local/global standby drives successful.

Changes made to Optimization will NOT take effect until all the logical drives are deleted and then the controller is RESET.

Click View to select one file.

Close transport fail at set configuration.

Collect mail canceled by user.

Communication is reestablished with controller.

Configuration information saved successfully!

Connect and Login command error

Contact Fax number format is wrong.

Contact Name is empty

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Contact Phone number format is wrong.

Continue add drv on logical drive *LogicalDrive*.

Controller Event, Battery %s %s. Informational message.

Controller Event, Battery %s %s. Likely battery module error. If error persists, replace defective battery module.

Controller Event, Controller Init Complete. Controller has been rebooted. Informational message.

Controller Event, Controller Reset. Informational Message.

Controller Event, DRAM Parity Error. Likely controller error. If error persists, replace defective controller.

Controller Event, Fan Back On-line. Informational message.

Controller Event, Fan Failure Detected. Fan is integrated with power supply. Replace power supply. Check LED indicators to identify supply.

Controller Event, Power Supply Back On-line. Informational message.

Controller Event, Power Supply Failure Detected. Replace power supply. Check LED indicators to identify supply.

Controller Event, Redundant Controller Failure Detected. Likely controller error. If error persists, replace defective controller.

Controller Event, SDRAM Error. Likely controller error. If error persists, replace defective controller.

Controller Event, Temperature Alert. Check room temperature, clogged or restricted airflow. Correct as necessary.

Controller Event, Temperature Back to Normal. Informational message.

Controller Event, UPS AC Power Back On. Informational message.

Controller Event, UPS AC Power Loss Detected. Contact technical support.

Controller has been reset.

Controllers back in redundancy mode!

Controllers back in redundancy mode!

Copy & Replace has completed on Physical drive %d:%d.

Copy & Replace has started from Physical drive %d:%d to %d:%d.

Copy & Replace on Physical drive %d:%d has been aborted.

Create logical drive failure.

Create logical volume failure.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Creating Host LUN filter Entry; please wait...

Creating standard Host LUN mapping; please wait...

Customer Fax number format is wrong.

Customer Name is empty.

Customer Phone number format is wrong.

Data is at risk. Media performance of this tape is severely degraded. Copy any data you require from this tape. Do not use this tape again. Restart the operation with a different tape.

Data Under/Overrun(Ch:%d,Id:%d). Likely poorly seated or defective drive. If random drives, possible I/O module or cable failure on drive loop.

Decrypt mail canceled by user.

Delete logical drive failure.

Delete logical volume failure.

Deleted disk successfully!

Deleting disk; please wait...

Deleting Host LUN filter Entry; please wait...

Deleting standard Host LUN mapping; please wait...

Device at [Channel:Target] has failed the self-reliability test.

Device information failed.

Disk is labeled successfully.

Door lock has an unknown status.

Door locked.

Door unlocked.

Download Firmware with Boot Record.

Downloading Firmware to the devices

Downloading Firmware to the RAID Controller.

Drive SCSI Ch:%d; Id:%d Likely poorly seated or defective drive. If random drives, possible I/O module or cable failure.

Email address format is wrong.

Enclosure #Enclosure state change: Name

Enclosure #Enclosure, temperature threshold has changed from a *State* state to a *State* state.

Engaging firmware; a controller reset is not necessary.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Error in writing file, please try later!

Error occurred. Reset the config file.

Expand LD/LV information.

Expand logical volume successfully.

Expansion has completed on logical drive *LogicalDrive*.

Expansion has started on logical drive *LogicalDrive*.

Expansion logical drive *LogicalDrive* has been aborted.

Fan #*Fan* has an unknown status.

Fan #*Fan* is malfunctioning.

Fan #*Fan* is not present in the system.

Fan #*Fan* is operational.

File I/O error. Configuration could not be restored!

Firmware downloaded to SAF-TE/SES device (Ch:%d,Id:%d) failed.

Firmware downloaded to SAF-TE/SES device (Ch:%d,Id:%d) successfully.

Firmware downloaded to SCSI drive (Ch:%d,Id:%d) failed.

Firmware downloaded to SCSI drive (Ch:%d,Id:%d) successfully.

First select one Logical Drive!

First select one physical device.

First select one RAID controller.

For RAID1 Logical Drive 'Add Drive'/'Copy and Replace' function is disabled.

Forward mail canceled by user.

Generic File Mutex lock is not released.

Get logical drive list failure.

Getting Agent Options failed!

Getting Controller Parameters failed!

Getting Host LUN Assignments failed! Please try again

Getting the configuration failed!

Gross Phase/Signal Error Detected (Ch:%d). Check drive seating, cables, and I/O module. Replace as required. Redundant controllers failed to bind.

Group Name can not be empty

Host LUN filter Entry created successfully!

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Host LUN filter Entry deleted successfully!

Host LUNs modified successfully.

If file access is not properly coordinated when assigned to multiple hosts, data corruption and access contentions may occur.

If Minimum interval is `\“0\”` or `\“*\”`, then Content must be `\“Event\”`

If the firmware download progress is interrupted the controllers/devices may become unusable.

Initialization has completed on logical drive *LogicalDrive*.

Initialization has started on logical drive *LogicalDrive*.

Initialization on logical drive *LogicalDrive* has been aborted.

Initialization operations have completed.

Initialization, rebuild, expand, or add SCSI drive activities are in progress. Try again later.

Inserted failed primary controller.

Inserted failed secondary controller.

Invalid data received.

Invalid server IP address!

IP Address format error (must be xxx.xxx.xxx.xxx and 0 < < xxx < < 255) !

IP Address is duplicated!

IP Address is empty.

Location is empty.

Log in as `\“ssconfig\”` user and try the operation again.

Logical array *LogicalArray* has changed from a *state* state to a *state* state.

Logical Drive *%d*, a SCSI Drive Failed. Replace defective drive.

Logical Drive ID *%d* exceeds 2 TB size limitation for sequential optimized LD.

Logical Drive ID *%d* exceeds 512 GB size limitation for random optimized LD.

Logical Drive ID *%d*, Rebuild Aborted. Informational message.

Logical Drive ID *%d*, Rebuild Completed. Informational message. Replace defective drive with new drive.

Logical Drive ID *LogicalDrive* rebuild has aborted.

Logical Drive ID *LogicalDrive* rebuild has completed.

Logical Drive ID *LogicalDrive* rebuild has started.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Logical Drive ID *LogicalDrive*, Add drive paused. Informational message.

Logical Drive ID *LogicalDrive*, Add SCSI Drv Aborted. Informational message.

Logical Drive ID *LogicalDrive*, Add SCSI Drv Completed. Informational message.

Logical Drive ID *LogicalDrive*, Add SCSI Drv Failed. Non-recoverable error. Likely drive failure or a non-recoverable error on stripe. Run parity check on the logical drive. Contact technical support.

Logical Drive ID *LogicalDrive*, Clone Aborted. Informational message.

Logical Drive ID *LogicalDrive*, Clone Completed. Informational message.

Logical Drive ID *LogicalDrive*, Clone Failed. Non-recoverable error. Likely drive failure or a non-recoverable error on stripe. Run parity check on the logical drive. Contact technical support.

Logical Drive ID *LogicalDrive*, Expand Aborted. Informational message.

Logical Drive ID *LogicalDrive*, Expand Completed. Informational message.

Logical Drive ID *LogicalDrive*, Expand Failed. Non-recoverable error. Likely drive failure or a non-recoverable error on stripe. Run parity check on the logical drive. Contact technical support.

Logical Drive ID *LogicalDrive*, Initialization Aborted. Informational message.

Logical Drive ID *LogicalDrive*, Initialization Completed. Informational message.

Logical Drive ID *LogicalDrive*, Initialization Failed. Informational message. Possible defective drive or slot. User may have aborted operation.

Logical Drive ID *LogicalDrive*, Parity check aborted by user. Parity check should be run to completion in order to verify data integrity of the specified Logical Drive.

Logical Drive ID *LogicalDrive*, Parity Check Completed. Informational message.

Logical Drive ID *LogicalDrive*, Rebuild Failed. Informational message. Possible defective drive or slot. User may have aborted operation.

Logical Drive ID *LogicalDrive*, Starting Add SCSI Drive Operation. Informational message.

Logical Drive ID *LogicalDrive*, Starting Clone. Informational message.

Logical Drive ID *LogicalDrive*, Starting Expansion. Informational message.

Logical Drive ID *LogicalDrive*, Starting Initialization. Informational message.

Logical Drive ID *LogicalDrive*, Starting Parity Check. Informational message.

Logical Drive ID *LogicalDrive*, Starting Parity Check. Informational message.

Logical Drive ID *LogicalDrive*, Starting Rebuild. Informational message.

Logical Drive *LogicalDrive* has changed from a *State* state to a *State* state.

Logical Drive *LogicalDrive* has changed from a *State* state to a *State* state.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Logical Drive *LogicalDrive*, Parity Check Failed. Non-recoverable error. Likely drive failure or a non-recoverable error on stripe. Run parity check on the logical drive. Contact technical support.

Note: If you manually aborted Parity Check, there is no risk of data corruption. There is no need to contact technical support.

Logical Volume Components List. Total disks in this Logical Volume=

Login successful

Login unsuccessful

Logout command error.

Mail information is empty

Memory Allocation Error. Unable to load daemon.

Minimum interval between emails must be one of the following formats: \“*\” \“0\”
“\nn\” \“nn:mm\”

Minimum interval can't meet Content value.

Minimum interval format error.

Minimum interval format is HH[:MM].

Minimum interval is empty.

Monitor Stopped.

Multiple IP assignment mechanisms are not supported. Select only one mechanism.

Mute beeper.

Mute controller beeper failure.

Mute controller beeper.

Mutex Lock is not released.

Mutex Lock timeouts.

Mutex Lock(s) cannot be created!

Mutex Lock(s) change state failed!

Negotiation Error (Ch:%d,Id:%d). Check drive seating, cables, and I/O module. Replace as required.

New parity check schedule has been created.

No admin progress exists.

No Admin progress was found. All activity should be complete!

No array administration activity in progress!

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

No more events to report.

No new controller was found; you do not need to reboot the system.

No online server, cannot delete event log.

No rebuildable drive available.

Not Ready Error Reported by Target (Ch:%d,Id:%d). Likely poorly seated or defective drive. Possible defective drive slot.

Only the last partition of a LD/LV can be deleted.

Open transport fail at set configuration.

Open transport, Lock fail, the original lock holder's IP address is %s.

Open transport, LogIn fail.

Operation completed successfully.

Operation failed.

Operation in progress.

Operation on one of the servers failed.

Out of the scope!

Parity Check Aborted.

Parity check confirmation.

Parity check could not be started on logical drive *LogicalDrive*.

Parity check has completed on logical drive *LogicalDrive*.

Parity check has started on logical drive *LogicalDrive*.

Parity check on logical drive *LogicalDrive* has been aborted.

Parity check schedule has been removed.

Parity check schedule has been updated.

Parity Error (Ch:%d,Id:%d). Likely poorly seated or defective drive. If random drives, possible I/O module or cable failure.

Parity Error (SCSI Channel %d, SCSI ID %d, Target Lun %d). Possible failure of I/O module, cable, controller, or HBA.

Partition 0 cannot be deleted.

Partition table is full. New partitions cannot be added.

Password error!

Password is empty.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Port is invalid (port must be between 0 and 65535).

Port is invalid (port must be between 1270 and 1273).

Power supply #*Power* has an unknown status.

Power supply #*Power* is malfunctioning and has been switched off.

Power supply #*Power* is malfunctioning or disabled.

Power supply #*Power* is not present in the system.

Power supply #*Power* is operational but disabled.

Power supply #*Power* is operational.

Power supply #*Power* is present in the system.

Primary controller failed!

Probe Agent command error!

Progress Not Available.

RAID Controller firmware checksum failed - corrupt firmware data.

RAID Controller firmware download has failed.

RAID Controller firmware has been updated.

RAID Controller is back online.

Read config file error.

Rebuild on logical drive *LogicalDrive* has been aborted.

Rebuild or check operations have completed.

Redundant Array State Change: Controller Failed.

Redundant Array State Change: Controllers Back in Redundancy Mode.

Removed a failed primary controller.

Removed a failed secondary controller.

Reselect Timeout (SCSI Channel %d, SCSI ID %d, Target Lun %d). Possible failure of I/O module, cable, controller, or HBA.

Restoring configuration to the controller.

Retry the configuration operation after the update is complete.

SAF-TE/SES card *Channel:Target* firmware has been updated.

SAF-TE/SES card *Channel:Target* global status has been updated.

SAF-TE/SES card *Channel:Target* slot perform status has been updated.

SAF-TE/SES card *Channel:Target* slot status has been updated.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Save a Copy in server.

Saving configuration information; please wait...

Scan SCSI drive (%d.%d) has succeeded.

Scan SCSI Drive information.

Scan SCSI Drive Successful (Ch:%d,Id:%d). Informational message.

Scheduled parity check could not be started on logical drive *LogicalDrive*.

Scheduled parity check has completed on logical drive *LogicalDrive*.

Scheduled parity check has started on logical drive *LogicalDrive*.

Scheduled parity check skipped to next schedule due to incompleteness of previous check.

SCSI Aborted Command (Ch:%d,Id:%d). Informational message. Drive may be failing or defective slot.

SCSI Bus Device Reset (SCSI Channel %d, SCSI ID %d, Target Lun %d). Possible failure of I/O module, cable, or HBA.

SCSI Bus Reset (SCSI Channel %d, SCSI ID %d, Target LUN %d). Possible failure of I/O module, cable, or HBA.

SCSI Bus Reset Issued (Ch:%d,Id:%d). Check drive seating, cables, and I/O module. Replace as required.

SCSI Channel Failure - Fibre Loop Connection. Failure Check drive path. Possible cable or module failure.

SCSI Channel Failure - General. Check drive path.

SCSI Channel Failure - Redundant Loop/Path Failure. Check drive path. Possible cable or module failure.

SCSI Channel Failure (Ch:%d,Id:%d). Replace defective I/O or drive module.

SCSI Channel Restored - Fibre Loop Connection Restored. Informational message.

SCSI Channel Restored - Redundant Path Restored. Informational message.

Secondary controller failed!

Server Error.

Server failed the operation due to insufficient memory.

Server Name is empty.

Server Not managed!

Set controller parameter(s) successful.

Set or add configuration failed.

Slot #*Slot* has had a device *Name*.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

SMART Predictable Failure Detected (Ch:%d,Id:%d). Informational message. Drive may be failing.

SMTP From address format is wrong or empty.

SMTP From address format is wrong.

SMTP Server address is empty.

Specify how often you want the parity checked.

Standard Configuration options provide a group of basic default RAID sets whereby the drive size, quantity, and assignments are preconfigured.

Standard Configuration will replace the existing configuration.

Standard Host LUN mapping created successfully!

Standard Host LUN mapping deleted successfully!

Standard maps are available to all connected hosts in certain configurations.

Standby rebuild operation completed with an unknown error.

Standby rebuild operation in progress.

Startup state of the secondary controller.

State Change Mutex lock is not released.

Sun StorEdge Configuration Service monitor daemon has started.

Sun StorEdge Configuration Service server daemon has started.

Sun StorEdge Configuration Service startup is complete.

Sun StorEdge Configuration Service trap daemon has started.

System Administrator (ssconfig) has started irrevocable System Shutdown and Restart. SAVE all data and LOG OFF IMMEDIATELY.

System ID is empty.

Tape operation has stopped because an error occurred while reading or writing data, which the drive cannot correct.

TapeAlert notification. Device [*Channel:Target*].

Target HW Error Reported (Ch:%d,Id:%d). Likely poorly seated or defective drive. Possible defective drive slot.

Target Media Error Reported (Ch:%d,Id:%d). Likely poorly seated or defective drive. Possible defective drive slot.

Target Timed-out (Ch:%d,Id:%d). Check drive seating, cables, and I/O module. Replace as required.

Temperature change from *TemperatureC* to *TemperatureC*.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

The background command (s) has finished.

The background command(s) has been aborted.

The battery on (Ch:%d,Id:%d) is expired on %s. Please install a new battery. Battery Information is (%s).

The battery on (Ch:%d,Id:%d) is going to expire on %s (after %d days). Please install a new battery before the current battery expired. Battery Information is (%s).

The changer mechanism is having difficulty communicating with the tape drive. Turn the autoloader off then on and restart the operation. If the problem persists, call the tape drive supplier helpline.

The configuration has been updated.

The configuration was successfully updated. If new LDs/LVs were created, the server may need to be rebooted.

The Controller devices list changed.

The controller parameters have been updated.

The current user is ssconfig; you cannot log in again.

The daemons are not responding.

The device does not belong to the same HBA card. The operations cannot continue.

The drive letter (mounted point) cannot be assigned.

The encrypt key is empty.

The firmware does not support multiple IP assignment mechanisms.

The firmware download failed because you have tried to use the incorrect firmware for this tape drive. Obtain the correct firmware and try again.

The Host LUN filter map (StorEdge SN#%d LD %d Partition %d WWN:) has been created.

The Host LUN filter map (StorEdge SN#%d LD %d Partition %d WWN:) has been deleted.

The in-service date of the new battery on (Ch:%d,Id:%d) is set.

The IP Address cannot be empty.

The last cleaning cartridge used in the tape drive has worn out. Discard the worn out cleaning cartridge, wait for the current operation to finish, and then use a new cleaning cartridge.

The LD# is moved up after you delete the LD, so it may not be consistent with the LD# shown in the RS-232 terminal window.

The length of the Encrypt Key must be greater than 8 characters.

The Mail Server field cannot be empty.

The media has exceeded its specified life.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

The memory allocation Failed.

The mount point (drive letter) cannot be written into file (registry).

The new battery on (Ch:%d,Id:%d) is installed. Battery information is (%s).

The number of logical drives did not increase after creating a logical drive. (before:%d now:%d).

The operation failed because the autoloader door is open. Clear any obstructions from the autoloader door, eject the magazine, and then insert it again. If the fault does not clear, turn the autoloader off and then on again.

The operation failed because the tape in the drive has snapped. Discard the old tape. Restart the operation with a different tape.

The Selected Device list is empty!

The server already exists!

The Server Name cannot be empty or only contain any extra spaces.

The server you selected might have already been shut down.

The set configuration failed because there is another conflicting operation.

The set configuration Failed.

The tape cartridge is not data grade. Any data you backup to the tape is at risk. The drive has not been able to read the MRS stripes. Replace the cartridge with a data-grade tape.

The tape drive has a hardware fault, which is not read/write related. Turn the tape drive off and then on again and restart the operation. If the problem persists, call the tape drive supplier helpline.

The tape drive has a hardware fault. Eject the tape or magazine, reset the drive (hold the unload button down for 5 seconds), and restart the operation.

The tape drive has a problem with the SCSI interface. Check the cables and cable connections and restart the operation.

The tape drive is due for routine cleaning. Wait for the current operation to finish and then use a cleaning cartridge.

The tape drive is having problems reading data. No data has been lost, but there has been a reduction in the capacity of the tape.

The tape drive needs cleaning. If the operation has stopped, eject the tape and clean the drive. If the operation has not stopped, wait for it to finish and then clean the drive.

The tape in the drive is a cleaning cartridge. If you want to backup or restore, insert a data-grade tape.

The tape is damaged or the drive is faulty. The drive can no longer read data from the tape. Call the tape supplier helpline.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

The tape is from a faulty batch or the tape drive is faulty. Use a good tape to test the drive. If the problem persists, call the tape drive supplier helpline.

The tape operation failed. Eject and then reinsert the tape or magazine.

The WWN must be a Hex string that is less than 16 characters.

There are no available disks to configure.

There are no LDs/LVs to manage.

There are not enough available disks to add a new LD/LV or there are no available LDs to create a LV. LDs must be unmapped and partitions must be deleted prior to being used in LVs.

There is a map to more than one Host, therefore, the \“Map to Multiple Hosts”\ property cannot be unchecked. Remove the map to multiple hosts before changing this setting.

There is a problem with the autoloader mechanism. Loader mechanism has detected a hardware fault.

There is no spare or replacement drive. Rebuild cannot continue.

This may cause data loss.

This operation can take up to one minute.

This operation might take several seconds. Please wait...

This operation will add to the existing configuration.

This operation will erase any existing configuration and data.

Unable to allocate memory for array inquiry. Unable to monitor this array.

Unable to determine local host, terminating!

Unable to eject cartridge because the tape drive is in use. Wait until the operation is complete before ejecting the cartridge.

Unable to inquire RAID controller.

Unable to inquire SAF-TE device. SAF-TE monitoring for this enclosure disabled.

Unable to inquire tape device. This device will be removed.

Unable to issue Self-Reliability Query on device at [*Channel:Target*].

Unable to log events.

Unable to open bindings file -or- no bindings present. Check Bindings file.

Unable to open I/O transport layer

Unable to query TapeAlert log data. TapeAlert monitoring disabled for this device.

Unable to read license data.

Unable to read response from SMTP server.

TABLE G-34 Error/Status Messages (*Continued*)

Error and Status Messages

Unable to send e-mail to SMTP server.

Unable to toggle the speaker on the SAF-TE/SES managed enclosure.

Unable to validate recipient at SMTP server.

Unable to write license data.

Undefined Error(Ch:%d,Id:%d). Contact technical support.

Unexpected Disconnect(Ch:%d,Id:%d). Check drive seating, cables, and I/O module. Replace as required.

Unexpected Sense Data Received from Target(Ch:%d,Id:%d). Likely poorly seated or defective drive. Possible defective drive slot.

Unexpected Unit Attention Received from Target(Ch:%d,Id:%d). Likely poorly seated or defective drive. Possible defective drive slot.

UNMAPPED LUN, NOT AVAILABLE TO HOST

Username is empty.

Wait for mutex failed.

Write config file error.

Wrong Fax number format.

Wrong Phone number format.

You must first delete the host mapping for the partitions whose index number is greater before you can change this partition.

You must first delete the host mapping for this LD/LV before you can delete it.

You must first delete the host mapping of the last partition before you can add a new partition.

You must have superuser administrator privileges to run this program. exiting...

Installation and Program Prompts

TABLE G-35 contains a list of installation and program prompts.

TABLE G-35 Installation and Program Prompts

Installation and Program Prompts

An existing configuration for this package is found. Do you want to restore it?

Are you sure you want to continue?

Are you sure you want to delete the log file?

Are you sure you want to delete the trap?

Are you sure you want to overwrite the file?

Are you sure you want to rebuild the drive?

Are you sure you want to restore factory defaults?

Before starting to use `SUNWsscsConsole`, run `/opt/SUNWsscs/sscsconsole/config_sscon` to input the pathname of the default Web browser to access HTML help files.

Begin downloading firmware to devices. Do you want to continue?

Cannot find the Java Runtime Environment
\$JRE_MAJOR.\$JRE_MINOR or above! `SUNWsscs Console` needs JRE 1.2 or above to run. The Console will try to find it in `/usr/java`, `/usr/jre`, `/usr/local/java`, or `/usr/local/jre`. If you installed the JRE in a directory other than one mentioned above, make a link.

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

Click View to select one file.

Do you want to download firmware to the controller?

Do you want to download firmware with boot record to the controller?

Do you want to save the current event log file?

Do you want to set up the mail server?

Enter a Contact Name.

Enter a Customer Name.

Enter a Location.

TABLE G-35 Installation and Program Prompts (*Continued*)

Installation and Program Prompts

Enter a mail address.

Enter a System ID.

Enter a valid gateway address.

Enter a valid IP address.

Enter a valid network mask.

Enter SMTP server information.

Enter SMTP Server information.

Enter the login password.

Enter the ssconfig user password to login.

Install JRE 1.2.2, 1.3, or 1.4 before installing the console.

Issue Reset to the controller?

Mapped LUNs exist! These must be deleted before creating new LDs/LVs. If you proceed, this operation overwrites the present configuration. Any existing LDs/LVs are deleted and all data is lost! If you want to keep your present configuration, use the "\Add LDs/LVs to the Current Configuration\" option. Do you want to continue?

Please select a LD/LV!

Please select CH/ID.

Please select the Channel/ID of which you want to scan.

Please specify at least one SSCS agent in command-line or in ssdgrcli.cfg.

Remove the package @PKG_NAME@ and then try to install this package.

Select a LD/LV.

Select a Product ID before downloading firmware.

Select Boot Record File.

Select Configuration File.

Select either AM or PM.

Select Firmware File.

Select only ONE LD/LV!

Select only one row to edit it.

Select only one row to remove it.

TABLE G-35 Installation and Program Prompts (*Continued*)

Installation and Program Prompts

Select Server for Managing Controller.

Select the agent before clicking Remove.

Select the agent item before clicking Edit.

Select the Channel/ID of the drive you want to scan

Select the firmware file before downloading it

Select the Logical Drive(s)

Select the server item before clicking Edit.

Select the server item before clicking Remove.

Select the Starting Day.

Select Timeout(Ch:%d,Id:%d). Likely poorly seated or defective drive. If random drives, possible I/O module or cable failure.

Specify Map Information.

SUNWsscscsConsole needs JRE 1.2.2, 1.3, or 1.4 to run. The Console will try to find it in /usr/java, /usr/jre, /usr/local/java, or /usr/local/jre. If you installed the JRE in a directory other than one mentioned above, make a link.

The client receives events only when it is running. Are you sure you want to exit?

The Content of Periodic Trap cannot be "Event". Please select another one.

The controller needs to be reset for the new firmware to take effect. Do you want to reset the controller?

The encrypt key you entered includes extra space in the beginning or at the end, which can cause encryption errors. Do you want to continue?

This may cause data loss. Are you sure you want to move?

This operation overwrites the present configuration. Any existing LDs/LVs are deleted and all data is lost! If you want to keep your present configuration, use the "Add LDs/LVs to the Current Configuration" option. Do you want to continue?

This operation overwrites the present configuration. Any existing LDs/LVs are deleted and all data is lost! If you want to keep your present configuration, use the "Add LDs/LVs to the Current Configuration" option. Do you want to continue?

This operation will result in the loss of current controller settings and data on ALL of the logical drives. Are you sure you want to restore factory defaults?

TABLE G-35 Installation and Program Prompts (*Continued*)

Installation and Program Prompts

This operation will result in the loss of data on ALL of the modified partitions. Do you want to continue?

To find out the java version currently installed, type "java-version"

Trap information has been modified. Do you want to save it?

You selected "Event" as the Trap Type, therefore, you must select at least one Active trap event.

You selected "My SMTP server needs authorization," therefore, you must enter a Password.

You selected "My SMTP server needs authorization," therefore, you must enter a Username.

You selected "use encrypt," therefore, you must enter an encrypt key.

Glossary

The glossary lists acronyms and defines RAID terms found throughout the documentation. It also includes definitions of the operational states for disk drives and logical drives.

- active-active controllers** A pair of components, such as storage controllers in a failure-tolerant RAID array, that share a task or set of tasks when both are functioning normally. When one component of the pair fails, the other takes the entire load. Dual active controllers are connected to the same set of devices and provide a combination of higher I/O performance and greater failure tolerance than a single controller.
- ANSI** American National Standards Institute.
- ARP** Address Resolution Protocol.
- automatic rebuild** A process in which data is automatically reconstructed after a drive failure and written to a standby (spare) drive. An automatic rebuild also occurs when a new drive is installed manually in place of a failed drive. If the rebuild process is interrupted by a reset, use the Manual Rebuild command from the NAS appliance firmware to restart the rebuilding process.
- block striping** *See striping.*
- block striping with dedicated parity** (RAID 3) This technique breaks data into logical blocks, the size of a disk block, and then stripes these blocks across several drives. One drive is dedicated to parity. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining drives.
- caching** Allows data to be stored in a predesignated area of a disk or RAM (random access memory). Caching is used to speed up the operation of RAID arrays, disk drives, computers and servers, or other peripheral devices.

- capacity** The total number of physical drives available for data storage in a RAID array (logical drive). For example, if the capacity is N-1 and the total number of disk drives in a logical drives is six 36-Mbyte drives, the disk space available for storage is equal to five disk drives (5 x 36-Mbyte or 180 Mbyte).
- CH** Channel.
- channel** Any path used for the transfer of data and control information between storage devices and a storage controller or I/O adapter. Also refers to one SCSI bus on a disk array controller. Each disk array controller provides at least one channel.
- CIFS** Common Internet File System.
- CISPR** International Special Committee on Radio Interference.
- DHCP** Dynamic Host Configuration Protocol.
- directory** Storage in a NAS appliance file system is typically organized into directories for sharing.
- disk mirroring** *See* mirroring (RAID1).
- EMC** Electromagnetic compatibility.
- EMU** Event monitoring unit.
- Fabric** Fibre Channel network built around one or more switches.
- Fabric switch** A Fabric switch functions as a routing engine that actively directs data transfer from source to destination and arbitrates every connection. Bandwidth per node via a Fabric switch remains constant when more nodes are added, and a node on a switch port uses a data path of up to 100 Mbyte/sec to send or receive data.
- failover** A mode of operation for failure-tolerant arrays in which a component has failed and its function has been assumed by a redundant component.
- fault tolerance** The capacity to cope with internal hardware problems without interrupting the array's data availability, often by using backup systems brought online when a failure is detected. Many arrays provide fault tolerance by using RAID architecture to give protection against loss of data when a single disk drive fails. Using RAID 1 (mirroring), RAID 3 or RAID 5 (striping with parity), or RAID 1+0 (mirroring and striping) techniques, the array controller can reconstruct data from a failed drive and write it to a standby or replacement drive.
- fault-tolerant logical drive** A logical drive that provides protection of data in the event of a single drive failure by employing RAID 1, 1+0, 3, or 5.
- FC-AL** (Fibre Channel-Arbitrated Loop) FC-AL is implemented as either a loop or a Fabric. A loop can contain up to 126 nodes, accessible through only one or two servers.

Fibre Channel	A cost-effective gigabit communications link deployed across a wide range of hardware.
Fibre Channel HBAs	Fibre channel adapters of a host computer, server, or workstation.
Fibre hubs	An Arbitrated Loop Hub is a wiring concentrator. "Arbitrated" means that all nodes communicating over this Fibre loop are sharing a 100 Mbyte/sec segment. Whenever more devices are added to a single segment, the bandwidth available to each node is further divided. A loop configuration allows different devices in the loop to be configured in a token ring style. With a Fibre hub, a Fibre loop can be rearranged in a star-like configuration because the hub itself contains port bypass circuitry that forms an internal loop inside. Bypass circuits can automatically reconfigure the loop once a device is removed or added without disrupting the physical connection to other devices.
FRU	Field-replaceable unit.
GB	Gigabyte. 1,000,000,000 (one billion) bytes.
GBIC	(Gigabit Interface Converter) A hot-swappable input/output device that plugs into a Gigabit Ethernet port or Fibre Channel.
global spare	A spare drive that is available to all logical drives in an array. Spare drives can be part of automatic logical drive rebuild.
group	In NAS, a group is a predefined grouping of users that require common directory access and privileges. After a group is created, it can be granted access to one or more share. Then any user that is added to the group has the share access and privileges that have been defined for that group. In other management programs, a group is a data object that enables multiple servers to be contained under a single category. Groups are similar in concept to domains, and enable you to organize servers.
HBA	Host bus adapter.
hot spare	A drive in a RAID 1 or RAID 5 configuration that contains no data and acts as a standby in case another drive fails.
hot-swappable	The ability of a field-replaceable unit (FRU) to be removed and replaced while the NAS appliance remains powered on and operational.
ID	Identifier number.
IEC	International Electrotechnical Commission.
initialization	The process of writing a specific pattern to all data blocks on all drives in a logical drive. This process overwrites and destroys existing data on the disks and the logical drive. Initialization is required to make the entire logical drive consistent at the onset. Initialization ensures that any parity checks performed in the future are executed correctly.

JBOD	(Just a Bunch of Disks) A storage device that consist of drives with no controllers.
LAN	Local area network.
LD	Logical drive.
load balancing	A feature that assigns NAS clients to specific Ethernet data ports when both data ports are in the same subnet to distribute the communication load across both data ports.
logical drive	A section of disk storage space, also referred to as a LUN, that is presented to the host operating environment as a single physical drive. A logical drive might be located on one or more physical drives. Each NAS appliance controller can manage one to eight logical drives.
LUN	(logical unit number) The major and minor device numbers make up the logical unit numbering sequence for a particular device connected to a computer.
LUN mapping	The ability to change the virtual LUN as presented to the server from storage. This enables such benefits as the ability of a server to boot from the SAN without requiring of a local disk drive. Each server requires LUN 0 to boot.
LUN masking	The characteristic that enables an administrator to dynamically map an HBA to a specified LUN. This provides an individual server or multiple servers access to an individual drive or to multiple drives, and prohibits unwanted server access to the same drives.
LVD	A low-noise, low-power, and low-amplitude signaling technology that enables data communication between a supported server and storage devices. LVD signaling uses two wires to drive one signal over copper wire and requires a cable that is no longer than 25 meters (82 ft.).
management port	The NAS appliance 10/100 Ethernet port (LAN2). The management port is the only port that can be used to configure a NAS appliance.
MB	(megabyte) 1,000,000 bytes or characters of data.
mirroring (RAID 1)	Data written to one disk drive is simultaneously written to another disk drive. If one disk fails, the other disk can be used to run the array and reconstruct the failed disk. The primary advantage of disk mirroring is 100 percent data redundancy. Since the disk is mirrored, it does not matter if one of the disks fails. Both disks contain the same data at all times and either can act as the operational disk. Disk mirroring provides 100 percent redundancy but is expensive because each drive in the array is duplicated.

**multiple-block striping
with distributed**

parity	A RAID technique (RAID 5) that offers redundancy with the parity information distributed across all disks in the logical drive. Data and its parity are never stored on the same disk. In the event that a disk fails, the original data can be reconstructed using the parity information and the information on the remaining disks.
NAS	Network attached storage.
NDMP	Network Data Management Protocol.
NFS	Network file system.
NVRAM	(non-volatile random access memory) A memory unit equipped with a battery so that the data stays intact even after main power is switched off.
N port	A Fibre Channel port in a point-to-point or Fabric connection.
out-of-band	Refers to the connections and devices that are not in the data path.
parity check	A process whereby the integrity of the redundant data on fault-tolerant arrays (RAID 3 and 5) is checked. The parity checking procedure on a logical drive recalculates the parity of data stripes in each of the logical drive's RAID stripe sets and compares it with the stored parity. If a discrepancy is found, an error is reported and the new correct parity is substituted for the stored parity.
partner group	A pair of interconnected controller units. Expansion units interconnected to the pair of controller units can also be part of the partner group.
PID	Primary controller identifier number.
RAID	(redundant array of independent disks) An arrangement of two or more disk drives combined into a single virtual drive to provide more disk storage space, better performance and reliability, and redundant backup of data. Various combinations of these features are described by defined RAID levels. Arrays can support RAID 0, 1, 1+0, 3, and 5.
RAID Level	Various techniques using combinations of mirroring, striping, duplexing, and parity to implement a RAID array are called RAID levels. Each technique uses a distinct algorithm to offer a mix of performance, reliability and cost.
RARP	Reverse Address Resolution Protocol.
read policy	A storage device parameter that determines whether the storage device holding data in cache before storing it to disk. The ability to hold data in cache while it is being written to disk can increase storage device speed during sequential reads.
rebuild	The process of reconstructing the data that was on a disk before it failed. Rebuilding can be done only in arrays with data redundancy, such as RAID levels 1, 1+0, 3, and 5.

- rebuild priority** Rebuild priority enables the RAID controller to serve other I/O requests while rebuilding the logical drives. Priority ranges from low, which uses the controller's minimum resources to rebuild, to high, which uses the controller's maximum resources to complete the rebuilding process.
- SAN** (storage area networking) A high-speed, open-standard, scalable network of storage devices and servers providing accelerated data access.
- SCSI** (Small Computer Systems Interface) An industry standard for connecting disk and tape devices to a workstation.
- SES** An interface to SCSI Enclosure Services devices. These devices sense and monitor the physical conditions within an enclosure, as well as enable access to the status reporting and configuration features of the enclosure (such as indicator LEDs on the enclosure).
- share** A share functions as a portal into a directory on a NAS appliance. Each share includes a set of permissions that apply to any group or user that accesses a directory through that share. Users must have access to a share to access files in a NAS directory.
- SID** Secondary controller identifier number.
- SMART** (Self-Monitoring Analysis and Reporting Technology) The industry-standard reliability prediction indicator for both the IDE/ATA and SCSI hard disk drives. Hard disk drives with SMART offer early warning of some hard disk failures so critical data can be protected.
- SMTP** (Simple Mail Transfer Protocol) A protocol for sending email messages between servers and from mail clients to mail servers. The messages can then be retrieved with an email client using either POP or IMAP.
- SNMP** (Simple Network Management Protocol) A set of protocols for managing complex networks. SNMP works by sending messages, called protocol data units (PDUs), to different parts of a network. SNMP-compliant devices, called agents, store data about themselves in Management Information Bases (MIBs) and return this data to the SNMP requesters.
- spanning** Making use of the firmware's striping capability to stripe data across two otherwise independent RAID logical drives. The two spanned logical drives are presented to the operating environment as one logical drive.
- standby drive** A drive that is marked as a spare to support automatic data rebuilding after a physical drive associated with a logical drive fails. For a standby drive to take the place of another drive, it must be at least equal in size to the failed drive and all of the logical drives dependent on the failed disk must be redundant—RAID 1, 1+0, 3, and 5.
- state** The current operational status of a disk drive, a logical drive, or controller. The NAS appliance stores the states of drives, logical drives, and the controller in its nonvolatile memory. This information is retained across power interruptions.

stripe size	<p>This is the amount of data in kilobytes that is striped across each physical drive in a logical drive. The values are in increments of 8 kilobytes and range from 8 to 64 kilobytes. Generally, large stripe sizes are more effective for arrays with sequential reads.</p> <p>To change the stripe size on an existing drive, you need to back up your data, redefine the stripe size, reconfigure the storage, and restore all the data.</p>
striping	<p>The storing of sequential blocks of incoming data on all the different SCSI drives in a logical drive.</p> <p>This method of writing data increases the disk array throughput because multiple drives are working simultaneously, retrieving and storing. RAID 0, 1+0, 3, and 5 and all use striping.</p>
terminator	A part used to end a SCSI bus. Terminators prevent energy from reflecting back into a cable plant by absorbing the radio frequency signals.
UPS	Uninterruptible Power Supply.
volume	Also called a logical unit number or LUN, a volume is one or more drives that can be grouped into a unit for data storage.
write-back cache	A cache-writing strategy in which the array controller receives the data to be written to disk, stores it in the memory buffer, and immediately sends the host operating environment a signal that the write operation is complete, without waiting until the data is actually written to the disk drive. Within a short time, the controller, when not busy, writes the data to the disk drive.
write policy	A cache-writing strategy used to control write operations. The write policy options are CIFS write-back and write-through cache.
write-through cache	A cache-writing strategy in which the array controller writes the data to the disk drive before signaling the host operating environment that the process is complete. Write-through cache has lower write operation and throughput performance than write-back cache, but it is the safer strategy, with minimum risk of data loss on power failure.
WWN	(worldwide name) A globally unique, hard-coded and embedded number assigned by the manufacturer and registered under IEEE that is used to identify hardware.

Index

A

- Add Server Grouping tab, 16
- Add Server Properties parameters, 13
- Add Server window, 13
- adding
 - host LUN, 55
 - logical drive, 128
 - new devices
 - probe, 79
 - SCSI drives to logical drive, 143
 - servers, 12
- Administrative Tools group, 30
- agent
 - Agent Options Management window, 22, 189
 - configurable parameters, 22
 - operation, 74
 - options management, 93
 - Setting Agent Options window, 24
 - SNMP agent, 209
- Agent Options Management command, 93
- Alarm State, 89
 - resetting, 89
- alarms
 - silencing, 89, 167
- Array Admin Progress command, 111, 116, 117, 118
- Array Admin Progress window, 93
- Array Administration
 - Array Admin Progress command, 111, 116, 117, 118
 - menu, 118
 - viewing progress, 111
- Array Administration Controller Assignment, 19

- arrays
 - maximum number supported, 19, 69, 92
- assign controller
 - through console, 18
- Assign Server to Manage RAID Controller window, 20
- Auto Discovery
 - button, 14
 - options, 75
 - process, 74
- automatic rebuild, 115
- Available SCSI IDs list box, 155
- Available Servers list
 - editing, 169

B

- battery
 - expiration information, 90
 - in-service date
 - setting, 91
 - status
 - critical, 90
 - degraded, 90
- beeper
 - muting, 122, 167
- before you begin installing Sun StorEdge Configuration Service
 - procedures, 3
- browser requirements
 - managing storage through the web, 107

- C**
- cache
 - write-back, 159
 - write-through, 159
 - Cache tab, 157
 - capacity
 - expanding, 141
 - Change Channel Settings window, 155
 - Change Controller Parameters, 153
 - Cache tab
 - Optimization, 159
 - Change Host LUN Assignments, 55
 - Change Settings button, 155
 - Channel Mode list box, 155
 - Clear button
 - configuration, 53
 - cluster configuration
 - CLUSTER.TXT, 201
 - IP addresses, 201
 - Microsoft Cluster Server software, 198
 - planning, 197
 - requirements, 198
 - setup, 198
 - start Sun StorEdge Configuration Service, 198, 200
 - Commit Changes Now, 30
 - Community String
 - checking, 215
 - configuration
 - caution, 35
 - Change Host LUN Assignments window, 55
 - clear button, 53
 - clearing, 53
 - Confirm Configuration Operation window, 39
 - Custom Configuration, 40
 - Custom Configure, 41
 - Custom Configure command, 153, 168
 - Load Configuration Confirmation window, 121
 - Load Configuration window, 120
 - loading, 119
 - logging out of ssconfig, 54
 - logical drive
 - saving, 27
 - Logical Drive Partitions box, 56
 - new
 - before you use, 42
 - New Configuration, 45
 - New Configuration option, 42
 - New Configuration warning, 45
 - NT Cluster option, 38
 - restoring, 119
 - Saveset Description tab, 120
 - saving, 27, 56
 - saving logical drive configuration, 56
 - Security Configuration window, 217
 - Security tab, 217
 - Select Configuration File window, 119
 - selecting RAID level, 43
 - setting traps, 216
 - Standard Configuration, 37
 - Standard Configuration window, 38
 - Standard Configure, 38
 - Use a standby drive check box, 38
 - verifying, 25
 - console
 - devices not displayed, 228
 - email setup, 212
 - running slowly, 229
 - controller
 - assigning managing server through the console, 19
 - converting from dual to single, 123
 - event errors, 243
 - muting, 122
 - parameters
 - changing, 153
 - resetting, 122, 154
 - selecting server to manage, 18
 - Controller Array Progress window, 117, 118
 - Controller Assignment command, 19
 - Controller Maintenance Options window, 122, 123, 147, 167
 - controller parameters
 - saving changes, 154
 - copying
 - member drives, 144
 - Custom Configuration, 40
 - Custom Configuration icon, 10
- D**
- deleting, 139
 - logical drive, 133
 - partition, 139

- devices
 - detailed view information, 76
 - not appearing in console, 228
 - rebuilding, 116
 - status, 71
- DHCP environment, 225
- Diagnostic Reporter
 - troubleshooting, 229
- Disk Access Latency field, 162
- Disk Administrator, 30
- Disk Array tab, 160
- Display HDD under LD check box, 11
- documentation
 - how book is organized, xvi
- Drive Check Period (Sec) field, 163
- Drive I/F tab, 161
- drive rebuild, 115, 116
- drives
 - copying and replacing, 144
 - make/change standby, 167
 - Physical Drive Standby State window, 168
 - rebuilding, 117
 - scanning in new, 146

E

- editing
 - available servers, 169
- email
 - Add Server, 214
 - Mail Address entry box, 214
 - Mail Server Setup window, 214
 - Mailing Lists, 213
 - Mailing Lists highlighted, 213
 - Mailing Lists tab, 14
 - messages, 212
 - Severity entry box, 214
 - SMTP email messages, 207
- Enable dynamic IP assignment, 166
- Enclosure Info tab, 82
- environmental
 - alarms not reported, 228
 - state, 88
- error codes, 231
 - administration errors, 241
 - client parameter errors, 238

- close transport errors, 239
- command code errors, 240
- communications async errors, 241
- communications link error, 240
- communications security errors, 241
- controller event errors, 243
- device state errors, 237
- disk state errors, 234
- drive side event errors, 244
- firmware download errors, 242
- host side event errors, 245
- initialization state errors, 238
- internal state errors, 237
- logical drive event errors, 245
- main communications errors, 240
- memory allocation errors, 239
- open transport errors, 239
- RDP path change errors, 243
- redundancy state errors, 236
- SAF-TE state errors, 235
- server manage/monitor event error, 245
- system drive state, 234
- system shutdown errors, 242
- tape state errors, 235
- timeout errors, 241
- transport field errors, 239

error messages, 246

- Event Log, 94
 - directory locations, 95
 - error codes, 231
 - Event Log file, 95
 - Event Log icon, 9
 - Event Log window, 95
 - eventlog.txt, 94, 95
 - Save Event Log icon, 9
 - severity levels, 98
- event messages
 - severity levels, 15
- expanding
 - capacity of logical drive, 141
 - partition, 141

F

- failed drive rebuild, 115
- firmware
 - upgrading

- controller, 147
- hard drives, 150
- SAF-TE devices, 152
- SES devices, 152

Format, 30

FRU

- definition, 92
- ID, obtaining, 92

G

- global spare drive
 - compared to a local spare drive, 186
 - definition, 174
- gray device status item, 71
- grouping servers, 16
- groups
 - colors, 73
 - Group Name, 12
 - status, 72
 - symbol colors, 72
 - View Group command, 77
 - View Group window, 77

H

- hard drives
 - displaying in the main window, 11
- HBA
 - dual port
 - main window, 78
- HBA device
 - adding manually, 62
- Host I/F tab, 163
- host ID
 - changing in a Fibre configuration, 156
- Host LUN
 - adding, 55
 - deleting, 56
- Host LUN Assignments box, 56
- HP OpenView, 207, 208
 - caution, 215
 - troubleshooting, 227
- HP-UX OS
 - creating logical volumes on servers, 32, 33

- event log location, 95
- SNMP traps, 219

I

- IBM AIX OS
 - SNMP traps, 220
 - updating the Object Data Manager, 171
 - writing event to log file, 96
- IBM NetView, 215
- icons
 - server with satellite dish attached, 25
 - status, 71
 - toolbar
 - Custom Configuration, 10
 - Event Log, 9
 - Save Event Log, 9
 - Save Report, 10
 - Server List Setup, 9
 - Standard Configuration, 10
- installing
 - Sun StorEdge Configuration Service on Solaris OS
 - before you begin procedures, 3
- IP address
 - changed in DHCP environment, 225
 - for traps, 218
- IP addresses
 - cluster configuration, 201
- Issue Controller Reset check box, 154, 165

J

- JBOD
 - discovering a drive
 - HP-UX OS, 194
 - IBM AIX OS, 194
 - Microsoft Windows OS, 193
 - Solaris OS, 192, 193
 - dual-bus configuration, 190
 - monitoring, 189, 197
 - single-bus configuration, 190
 - upgrading firmware on hard drives, 191

L

- label
 - automatically writing, 49
- Linux OS
 - event log location, 95
 - SNMP traps, 219
 - stopping console, 7
- loading configuration, *See* restoring logical drive configuration
- local spare drive
 - compared to a global spare drive, 185
 - definition, 174
- logging in, 17
- logging out, 17
- logical drive
 - adding, 128
 - adding SCSI drives, 143
 - creating a partition, 136
 - definition, 174
 - deleting, 133
 - deleting a partition, 139
 - drive allocation, 176
 - expanding capacity, 141
 - maximum number of disks per logical drive, 158
 - maximum number of drives for RAID level, 47
 - maximum usable capacity per logical drive, 158
 - Partitions box, 56
 - restoring configuration, 119
 - saving configuration, 27
 - View Logical Drive command, 86
 - View Logical Drive window, 88
- logical drive number, 135
- logical volume
 - creating, 51, 130
 - definition, 174
 - no logical drives listed, 225
- logical volumes
 - avoiding, 36
- LUN assignments
 - maximum, 54
 - maximum number, 36
- LUN filter
 - accessing LUN Filter view, 61
 - adding HBA device manually, 62
 - mapping logical drive to host, 65
 - overview, 59
 - removing standard mapping, 64

- LUN Filter Properties, 64
- LUNs, 29
 - adding, 55
 - Change Host LUN Assignments, 54
 - definition, 175
 - deleting, 56
 - Host LUN Assignments box, 56
 - LUN detailed information, 81
- LUNs Per Host parameter, 163

M

- main window, 69
- Major field, 232
- Managed Servers list setup, 17
- Managing storage through the web, 107
- mapping partition to LUN description, 178
- Max Queued IO Count field, 163
- member drives
 - copying and replacing, 144
- menu bar
 - overview, 8
- message variables, 246
- MIB, 209
- Microsoft Cluster Server (MSCS), 197
- Microsoft Windows OS
 - creating partitions, 31
 - determining worldwide name, 204
 - event log location, 95
 - starting Sun StorEdge Configuration Service, 7
 - Windows NT
 - using cluster configuration (SCSI only), 197
 - Windows NT partitions, 30
- Minor field, 233
- monitoring process, 74
- mountall command, 29
- muting
 - alarm, 167
 - beeper, 167

N

- navigating, 11
- network

- DHCP, 166
- RAAP, 166
- static IP, 166
- Network tab, 166
- newfs command, 29
- NT Cluster option, 38
- NTFS, 199

O

- Object Data Manager
 - updating on IBM AIX OS, 171
- online help, 71
- optimization
 - random I/O
 - maximum size, 158
 - stripe size, 157
 - sequential I/O
 - maximum size, 158
- out-of-band
 - storage management, 101
- Out-of-Band Agent preferred over In-Band check box, 104

P

- parameters
 - controller, 153
- parity check, 112
- scheduling, 113
- partitioning
 - creating, 136
 - Microsoft Windows, 30
 - Solaris, 30
 - deleting, 139
- partitioning, 56, 139
 - an existing logical drive, 136
 - requirements, 44
- performance statistics, 124
- physical devices
 - Physical Drives tab, 81
 - View Physical Device window, 87
 - View Physical Drive, 87
- Physical Drive Standby State window, 168
- physical drives

- displaying in the main window, 11
- polling
 - status change of devices, 23
- Power Up SCSI Reset field, 161
- probe, 79
- Product ID list box, 151
- prompts, 263
- purple device status icon
 - device status state, 71

R

- RAID
 - Enclosure Info tab window, 82
 - RAID controller parameters, 83, 86
 - terminology overview, 173
 - View Controller Configuration command, 80
- RAID levels
 - assigning to logical drive, 39
 - RAID 0, 180
 - RAID 1, 181
 - RAID 1+0, 182
 - RAID 3, 183
 - RAID 5, 184
 - range of disks/logical drives supported, 179
- random I/O optimization
 - maximum size, 158
 - stripe size, 157
- rebuild
 - drive manually, 117
 - failed drive, 115
 - standby drive (automatic), 115
- Rebuild Priority, 160
- Rebuild window, 118
- rebuilding
 - drive without standby, 116
- red device status item, 71
- Redundancy tab, 164
- Redundant Deassert Reset field, 164
- Redundant Force Pri Fail field, 164
- Redundant Force Sec Fail field, 164
- Report icon, 10
- report option, 98, 101
- Reset button, 228
- reset button

- alarm state box, 89
- silencing failed component alarms, 89
- View Enclosure window, 89
- resetting
 - controller, 122
- restoring
 - logical drive configuration, 119
- restoring a logical drive configuration, 119
- RS 232 tab, 156
- RSAGENT.DLL, 227
- RST_OID.MIB file, 208, 226

S

- SAF-TE Polling Period (Sec) field, 162
- Save Configuration function, 28, 57
- Save Configuration window, 28
- save controller parameter changes, 154
- save logical drive configuration, 56
- Save Report command, 98
- Save Report File window, 99
- Saveset Description tab, 120
- saving logical drive configuration, 27
- scanning
 - drives, 146
- scheduling
 - parity check, 113
- SCSI channel
 - definition, 175
- SCSI drives
 - adding to logical drive, 143
- SCSI I/O Timeout (Sec) field, 162
- SCSI Motor Spin Up field, 161
- Select Configuration File window, 28, 57
- Select Controller Progress window, 117
- sequential I/O optimization
 - maximum size, 158
- Server List Setup window, 169
- server not responding, 223
- Server Setup window, 170
- servers
 - Add Server Properties, 13
 - Add Server window, 13, 170, 213
 - add to Managed Servers list, 12
 - assign to controller through console, 18
 - Available Servers list, 170, 212
 - Edit Server window, 170, 213, 225
 - editing, 169
 - inventory, 74
 - Managed Servers list, 170, 225, 226
 - managing, 17
 - manually deleting managing server using
 - terminal, 22, 124
 - organize into groups, 16
 - Server List Setup command, 169
 - Server List Setup function, 201, 212
 - Server List Setup icon, 9
 - Server List Setup procedure, 12
 - Server List Setup window, 226
 - Server Login window, 18
 - Server Setup, 225
 - Server Setup window, 169, 212, 226
 - setup to send traps, 215
 - TCP/IP address for server, 13
 - View Server command, 78
 - View Server window, 78, 79
- Service Configuration window, 216
- Set Controller Config field, 164
- Severity field, 231
- severity levels (events), 98
- silencing
 - environmental alarms, 89
- Simple Mail Transport Protocol (SMTP), 214
- SMART Drive field, 88
- SMART monitoring
 - enabling, 24
- SMTP e-mail messages, 207
- SN#, 70
- SNMP
 - agents and managers, 209
 - how it works, 208
 - identifying objects, 209
 - Management Information Base, 209
 - objectIdentifier of MIB variable, 211
 - protocol data unit (PDU), 210
 - request types, 210
 - security, 210
 - software model, 209
 - V1, 210
- SNMP traps
 - messages description, 208

- setting
 - using Sun StorEdge Configuration Service, 212
 - without using Sun StorEdge Configuration Service, 221
- setup, 215
- Solaris OS
 - event log location, 95
 - LG not present, 228
 - rstrapd, 226
 - SNMP traps, 218
 - stopping console, 7
- ssmon password bypassed, 14
- sstrapd daemon, 218, 219, 220, 221
- Standard Configuration, 37
- Standard Configuration icon, 10
- standby drive
 - making/changing, 167
- starting
 - Sun StorEdge Configuration Service
 - Microsoft Windows OS, 7
 - Solaris, Linux, HP-UX, and IBM AIX OS, 6
- status
 - device, 71
 - groups, 72
- storage
 - managing through the web, 107
 - browser requirements, 107
- stripe size
 - random optimization, 157
 - sequential optimization, 157
- Sun StorEdge 3120 SCSI array
 - monitoring, 189
 - supported functionality, 4
- Sun StorEdge Configuration Service
 - agent operation, 74
 - console, 69
 - logging in/logging out, 17
 - main window, 69
 - monitoring process, 74
 - overview, 1
 - starting, 6
 - tree view, 72
- SuperFlex storage array, 199
- system user password usage, 14

T

- tabs
 - described, 11
- Tag Count Per drive field, 162
- TCP/IP
 - establishing connection, 74
- toolbar icons, 9
- trap daemon
 - troubleshooting, 227
- trap messages, 208
- traps
 - Card Name, 208
 - Community Name box, 217
 - Event Date, 208
 - Event Severity, 208
 - Event Time, 208
 - Internet Protocol (IP), 208
 - IP address format, 218
 - Microsoft SNMP Properties window, 217
 - not received from server, 226
 - OID (object identifier), 208
 - on HP-UX OS, 219
 - on IBM AIX, 220
 - on IBM AIX OS, 220
 - on Linux, 219
 - on Solaris, 218
 - Server Address and Name, 208
 - Service Configuration window, 217
 - SNMP trap messages, 208
 - Transmission Control Protocol (TCP), 208
 - Trap Destinations list, 217
 - trap recipient setup, 217
 - Traps tab, 217
- traps setup, 215
- tree view, 72
- troubleshooting, 223
 - See also* error codes
 - See also* error messages

U

- unmapping
 - LUNs, 56
- upgrading
 - firmware
 - controller, 147

- hard drives, 150
- SAF-TE devices, 152
- SES devices, 152
- Use a standby drive check box, 38

V

- verifying configuration, 25
- View CD-ROM window, 87, 88
- View Controller Configuration, 80
- View Enclosure command, 88
- View Enclosure window, 88, 89
- View FRU, 92
- View FRU window, 92
- View HBA Card, 79
- View menu, detailed information, 76
- View Report command, 101
- View Tape window, 88

W

- warning
 - New Configuration, 45
- white device status item, 71
- worldwide name
 - determining on Linux OS, 204
 - determining on Microsoft Windows OS, 204
 - determining on Solaris OS, 203
- worldwide node name
 - determining, 205
- worldwide port name
 - determining, 206
- Write a new label to the new LD check box, 49, 121
- write-back cache
 - enabling, 159
- write-through cache
 - enabling, 159

Y

- yellow device status item, 71

