



Sun StorEdge 6920 System Administration Guide

For the Browser Interface Management Software

Sun Microsystems, Inc.
www.sun.com

Part No. 819-0123-10 (V3)
September 2005

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

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

Use is subject to license terms.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo, Java, Jiro, Netra, Solaris, and Sun StorEdge are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Legato and the Legato logo are registered trademarks, and Legato NetWorker are trademarks or registered trademarks of Legato Systems, Inc.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

Use of any spare or replacement CPUs is limited to repair or one-for-one replacement of CPUs in products exported in compliance with U.S. export laws. Use of CPUs as product upgrades unless authorized by the U.S. Government is strictly prohibited.

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 © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS L'AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

L'utilisation est soumise aux termes de la Licence.

Cette distribution peut comprendre des composants développés par des tierces parties.

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

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

L'utilisation de pièces détachées ou d'unités centrales de remplacement est limitée aux réparations ou à l'échange standard d'unités centrales pour les produits exportés, conformément à la législation américaine en matière d'exportation. Sauf autorisation par les autorités des Etats-Unis, l'utilisation d'unités centrales pour procéder à des mises à jour de produits est rigoureusement interdite.

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



Please
Recycle



Adobe PostScript

Contents

1. Product Overview 1

System Overview 2

About the Sun StorEdge 6920 System 2

About High-Availability Features 4

About the Sun StorEdge 6920 Hardware 4

About the System Cabinet 5

About the Service Processor Panel 6

About the Storage Service Processor 8

About the Storage Service Processor Accessory Tray 9

About the Data Services Platform 9

About Storage Arrays 12

About the Sun StorEdge 6920 Software 14

About User Interfaces 14

About Resident Software 14

About Required Host Software 15

About Data Services 17

About User Roles 18

Configuration Overview 19

About the Default State of the System 19

- About Provisioning Storage 21
- About Physical Storage Elements 24
- About Logical Storage Elements 25
- Getting Started With Configuration 26
- Monitoring Overview 28
 - How the Monitoring Software Works 28
 - About Monitoring Strategy 29
 - About the Monitoring Cycle 29
 - About the Event Life-Cycle 30
- Getting Started With Monitoring 32
- 2. Using the Browser Interface 33**
 - About the Browser Interface 34
 - About the Page Layout 34
 - About the Page Banner 34
 - About the Navigational Tabs 36
 - About the Page Content Area 38
 - Controlling the Display of Table Information 38
 - Using Forms 41
 - Searching for System Elements 41
 - Using Help 42
 - Logging Out of the Management Software 43
- 3. Performing Connection and Administration Tasks 45**
 - Connecting the System 46
 - About Network Domains 46
 - About LAN Names and IP Addresses 47
 - About Internal System IP Addresses 48
 - About Network Configuration Options 50

Managing From the Site LAN	51
Setting Up Connections to the Site LAN	51
About Data Host Connections	51
About Restricting Services	52
About Allocating Ports	53
About Storage Resource Card (SRC) Sets	54
Guidelines for Fibre Channel Port Connections	55
About Host Bus Adapters	55
Connecting an External Switch	56
About Multipathing	57
Administering a System Isolated From the LAN	57
Installing the Remote Scripting CLI Client	58
About Lights-Out Management	58
Powering the System On and Off	59
Performing a Partial Shutdown	59
Restoring the System After a Partial Shutdown	60
Performing a Full Shutdown	60
Restoring the System After a Full Shutdown	61
Specifying General Administrative Settings	63
Adding a System Description	63
Setting the System IP Address	63
Enabling the NTP Server	64
Setting the System Time	65
Changing Passwords	65
Performing Monitoring Administration	67
General Setup Tasks	67
Displaying and Editing Site Setup Information	67
Displaying and Editing System Timeout Settings	68

Displaying and Editing System Configuration Options	68
Enabling/Disabling Fault Signature Analysis	69
Configuring Parallel Monitoring Parameters	69
Displaying and Editing MAC Addresses (Ethers)	70
Exporting a Configuration	70
About Notification	71
Email Configuration Tasks	71
Configuring Email Notification	72
Testing a Notification Email Address	72
Adding an Email Notification Recipient	73
Editing an Email Notification Recipient	73
Deleting an Email Notification Recipient	74
Using Email Filters	74
About Email Filters	74
Adding an Email Filter	75
Editing an Email Filter	75
Deleting an Email Filter	75
Remote Notification Tasks	76
Configuring Remote Notification	76
Adding an SNMP Notification Recipient	76
Editing an SNMP Notification Recipient	77
Deleting an SNMP Notification Recipient	77
Editing Notification Setup Parameters	78
Agent Tasks	78
Activating/Deactivating the Agent	78
Running the Agent	79
Displaying Agent Statistics	79
Clearing Agent Statistics	79

Ejecting a CD-ROM 80

4. Performing Configuration and Management Tasks 81

Configuring Feature Licensing 82

 About Feature Licenses 82

 Displaying Feature License Compliance 84

 Adding a Feature License Key 84

 Displaying a Feature License Key 85

 Deleting a Feature License Key 85

Configuring Storage Volumes 87

 About Volumes 87

 Planning Volumes 88

 Managing Volumes 88

 Displaying Volume Information 89

 Creating a Volume 89

 Modifying a Volume 90

 Deleting a Volume 90

 Mapping a Volume to Initiators 90

 Unmapping a Volume From Initiators 91

 Extending a Volume 91

 Extending Replicated Volumes 92

Configuring Mirrored Volumes 93

 About Mirroring 93

 Planning Mirror Components 95

 Managing Mirror Components 97

 Displaying Mirror Component Information 97

 Creating a Mirrored Volume 97

 Mirroring an Existing Volume 98

 Adding a Mirror Component 99

Splitting a Mirror Component	99
Rejoining a Split Mirror Component	100
Reverse Rejoining a Split Mirror Component	101
Breaking a Mirror Component	101
Forcing the Break of a Mirror Component	102
Configuring Snapshots	103
About Snapshots	103
About Snapshot Reserve Space	104
About Snapshot Rollbacks	105
Planning Snapshots and Snapshot Reserve Space	106
Managing Snapshot Reserve Space	108
Allocating Snapshot Reserve Space	108
Expanding Snapshot Reserve Space	109
Deleting Snapshot Reserve Space	110
Managing Snapshots	111
Displaying Snapshot Information	111
Selecting the Snapshot Policy	111
Creating a Snapshot	112
Resnapping a Snapshot	113
Deleting Snapshots	113
Rolling Back a Snapshot	114
Canceling a Rollback Job	114
Using Data Replication	116
About Data Replication	116
About the Replication Process	116
About Replication Sets	118
About Consistency Groups	121
About Replication Links	123

About Replication Set Properties	126
About Primary and Secondary Role Reversal	128
About Replication Modes	129
Reference: Synchronous Versus Asynchronous Replication Modes	131
About Data Replication Operations	132
Planning for Data Replication	134
General Planning Considerations	135
Reference: Data Replication Tasks	136
Planning to Create Replication Sets and Consistency Groups	137
Managing Data Replication	139
Displaying Replication Link Information	139
Displaying Replication Set Information	140
Displaying Consistency Group Information	140
Creating Replication Sets and Consistency Groups	141
Deleting Replication Sets	141
Adding a Replication Set to a Consistency Group	142
Combining Replication Sets in a Consistency Group	143
Removing a Replication Set From a Consistency Group	144
Deleting Consistency Groups	145
Reversing Roles	145
Changing Replication Modes	146
Replicating Data	147
Synchronizing Data Using a Backup Tape	148
Troubleshooting Data Replication	148
Configuring Virtual Disks	150
About Virtual Disks	150
Displaying Virtual Disk Information	151
Deleting a Virtual Disk	151

Configuring Storage Pools	153
About Storage Pools	153
Planning Storage Pools	153
Managing Storage Pools	154
Displaying Pool Information	154
Creating a Storage Pool	155
Adding Storage to a Pool	155
Modifying a Storage Pool	156
Deleting a Storage Pool	156
Configuring Storage Profiles	158
About Storage Profiles	158
About Virtualization Strategies	161
Reference: Standard Storage Profiles	162
Managing Storage Profiles	164
Displaying Profile Information	164
Creating a Storage Profile	165
Modifying a Storage Profile	165
Deleting a Storage Profile	166
Configuring Storage Domains	167
About Storage Domains	167
Managing Storage Domains	169
Displaying Storage Domain Information	169
Creating a Storage Domain	169
Deleting a Storage Domain	170
Configuring Initiators	171
About Initiators	171
Managing Initiators	171
Displaying Initiator Information	172

Creating an Initiator	172
Mapping an Initiator to Volumes	173
Deleting an Initiator	173
Unmapping an Initiator From a Volume	174
Configuring System Ports	175
About Ports	175
Displaying Port Information	176
Displaying Replication Link Information	177
Enabling Ports for Data Replication	177
Disabling Ports for Data Replication	178
Adding Replication Links	179
Deleting Replication Links	179
Configuring Storage Arrays	180
About Arrays and Trays	180
About Conventions for Array Names and IP Addresses	180
Displaying Storage Array Information	182
Modifying an Array	183
Using Storage Trays	184
About Storage Trays	184
About Tray Numbering Conventions	185
About Identifiers for Cabinets	186
Displaying Storage Tray Information	188
Managing Disk Drives	189
About Disk Drives	189
Displaying Disk Information	190
Managing External Storage	191
About Using External Storage	191
Planning to Use External Storage	194

- Displaying External Storage Information 199
- Creating a Legacy Volume 200
- Migrating Externally Stored Data 201
- Using External Storage as Raw Storage 202

5. Monitoring the System 205

- Setting Diagnostic Viewing Preferences 206
- Monitoring Alarms and Events 207
 - About Alarms and Events 207
 - About Alarm Management 207
 - Displaying Alarm Information 208
 - Displaying an Alarm Summary 208
 - Displaying Alarm Details 208
 - Displaying Aggregated Events 208
 - Displaying Alarms for a Specific Device 209
 - Displaying an Alarm Statistics Report 209
 - Acknowledging Alarms 210
 - Reopening Acknowledged Alarms 210
 - Deleting Alarms 211
 - Generating Alarm Statistics Reports 211
 - Displaying Event Information 212
 - Displaying the Event Log 212
 - Displaying Event Details 212
- Monitoring Devices 213
 - Displaying Device Information 213
 - Displaying a Rack Inventory Summary 213
 - Displaying Rack Components 214
 - Displaying Device Details 214

Displaying Device FRUs, Components, or Logical Components of a Specific Type	214
Displaying Detailed Information for a Specific Device FRU, Component, or Logical Component	215
Displaying System Parameters for a Device	216
Displaying Performance Data	216
Displaying Rack Performance Data	216
Displaying Performance Data for a DSP or Array	217
Displaying DSP Utilization or Replication Queue Depth	217
Adding Notes for a Rack	218
Adding Notes for a Device	218
Generating New Inventory	219
Changing the Monitoring Status for a Device	219
Setting Up Volume Verify	220
Installing and Uninstalling Revision Patches	221
About Installing and Uninstalling Revision Patches	221
Generating a Revision Maintenance Upgrade Report	222
Installing Revision Patches	222
Generating a Revision Maintenance Backout Report	223
Backing Out Revision Patches	224
Monitoring Topologies	225
Displaying the System Topology	225
Changing Topology Views	225
Archiving a Topology	226
Displaying Archived Topologies	226
Displaying an Archived Topology Snapshot	226
Comparing Two Archived Topology Snapshots	227
Deleting an Archived Topology Snapshot	227
Managing Storage Configuration Jobs	228

- About Configuration Jobs 228
- Reference: Configuration Job Triggers 228
- Displaying Configuration Job Information 230
- Canceling and Removing Configuration Jobs 231
- Managing Storage Monitoring Jobs 232
 - About Monitoring Jobs 232
 - Displaying Monitoring Job Information 232
 - Displaying a Summary of Monitoring Jobs 232
 - Displaying the Details of a Monitoring Job 233
 - Displaying a Summary of Archived Monitoring Jobs 233
 - Displaying the Details of an Archived Monitoring Job 233
 - Stopping a Monitoring Job 234
 - Archiving Monitoring Jobs 234
 - Deleting Monitoring Jobs 234
 - Deleting Archived Monitoring Jobs 235

6. Servicing and Troubleshooting the System 237

- Keeping the System Up to Date 238
- Servicing Hardware Components 239
- About Field-Replaceable Units 240
- Troubleshooting LEDs 241
 - Interpreting LED Status Indicators 241
 - Interpreting Data Services Platform (DSP) LEDs 242
 - Interpreting Card Status LEDs 245
- Troubleshooting With the Monitoring Software 256
 - Testing an Email Address 256
 - Generating a Solution Extract 257
 - Displaying a Solution Extract 257
 - Working With Core Files 258

About Core Files 258

Downloading and Deleting Core Files 258

Running a Diagnostic Test on a Device 259

Responding to System Error Messages 260

Viewing Log Events 271

Glossary 273

Index 289

Preface

The *Sun StorEdge 6920 System Web-Based Online Help* provides system overview, configuration, maintenance, and basic troubleshooting information about the Sun StorEdge™ 6920 system. This document is the printable version of the online help, and as such, it is not organized as a typical printed document. The intent of this document is to provide early access to the Sun StorEdge 6920 system online help so that you can better plan for the system.

Before You Read This Book

Before you read this book, it is recommended that you prepare the installation site and install and configure the Sun StorEdge 6920 system as described in these books:

- *Sun StorEdge 6920 System Regulatory and Safety Compliance Manual*
- *Sun StorEdge 6920 System Site Preparation Guide*
- *Sun StorEdge 6920 System Getting Started Guide*

How This Book Is Organized

Chapter 1 introduces the product.

Chapter 2 describes the browser interface and how to use it.

Chapter 3 provides information about connecting the system and performing administrative tasks.

Chapter 4 provides information about system configuration and management tasks.

Chapter 5 describes how to monitor the system.

Chapter 6 provides servicing and troubleshooting information.

Sun StorEdge 6920 System Documentation

The following is a list of documents related to the Sun StorEdge 6920 system. For any document number with *nn* as a suffix, use the most current document.

Subject	Title	Part Number
Unpacking instructions attached to the shipping container	<i>Unpacking Guide</i>	816-6385- <i>nn</i>
Data services license information	<i>Sun StorEdge Storage Pool Manager Software for Non-Sun Systems, Version 1.0, Right to Use License</i>	819-2404- <i>nn</i>
	<i>Sun StorEdge Data Mirroring Software, Version 1.0, Right to Use License</i>	819-2335- <i>nn</i>
	<i>Sun StorEdge Data Replicator-Synchronous Software, Version 1.0, Right to Use License</i>	819-2336- <i>nn</i>
	<i>Sun StorEdge Data Replicator-Asynchronous Software, Version 1.0, Right to Use License</i>	819-2337- <i>nn</i>
System planning information	<i>Sun StorEdge 6920 System Site Preparation Guide</i>	819-0118- <i>nn</i>
System regulatory and safety information	<i>Sun StorEdge 6920 System Regulatory and Safety Compliance Manual</i>	819-0119- <i>nn</i>
Installation and configuration information	<i>Sun StorEdge 6920 System Getting Started Guide</i>	819-0117- <i>nn</i>
Late-breaking information not included in the information set	<i>Sun StorEdge 6920 System Release Notes</i>	819-0120- <i>nn</i>

The `sscs(1M)` man page provides information about the commands used to manage storage using the command-line interface (CLI).

Related Documentation

The following documents provide detailed information, if needed, about related system components.

Product	Title	Part Number
Best practices	<i>Best Practices for Sun StorEdge 6920 System (Version 3.0.0)</i>	819-0122- <i>nn</i>
Sun Storage Automated Diagnostic Environment Enterprise Edition	<i>Sun Storage Automated Diagnostic Environment Enterprise Edition Release Notes Version 2.4</i>	819-0432- <i>nn</i>
SAN Foundation software	<i>Sun StorEdge SAN Foundation 4.4 Configuration Guide</i>	817-3672- <i>nn</i>
Traffic Manager software	<i>Sun StorEdge Traffic Manager 4.4 Software Release Notes For HP-UX, IBM AIX, Microsoft Windows 2000 and 2003, and Red Hat Enterprise Linux</i>	817-6275- <i>nn</i>
	<i>Sun StorEdge Traffic Manager 4.4 Software User's Guide For IBM AIX, HP-UX, Microsoft Windows 2000 and 2003, and Red Hat Enterprise Linux</i>	817-6270- <i>nn</i>
	<i>Sun StorEdge Traffic Manager 4.4 Software Installation Guide For Red Hat Enterprise Linux</i>	817-6271- <i>nn</i>
	<i>Sun StorEdge Traffic Manager 4.4 Software Installation Guide For Microsoft Windows 2000 and 2003</i>	817-6272- <i>nn</i>
	<i>Sun StorEdge Traffic Manager 4.4 Software Installation Guide For IBM AIX</i>	817-6273- <i>nn</i>
	<i>Sun StorEdge Traffic Manager 4.4 Software Installation Guide For HP-UX 11.0 and 11i</i>	817-6274- <i>nn</i>
Sun StorEdge network Fibre Channel switch-8 and switch-16	<i>Sun StorEdge Network 2Gb FC Switch-8 and Switch-16 FRU Installation</i>	817-0064- <i>nn</i>
	<i>Sun StorEdge Network 2 Gb FC Switch-8 and Switch-16 Release Notes</i>	817-0770- <i>nn</i>
	<i>Sun StorEdge Network 2 Gb FC Switch-64 Release Notes</i>	817-0977- <i>nn</i>
Sun StorEdge Brocade switch documentation	<i>Sun StorEdge Network 2 Gb Brocade SilkWorm 3200, 3800, and 12000 Switch 3.1/4.1 Firmware Guide to Documentation</i>	817-0062- <i>nn</i>
Sun StorEdge McData switch documentation	<i>Sun StorEdge Network 2 Gb McDATA Intrepid 6064 Director Guide to Documentation, Including Firmware 5.01.00</i>	817-0063- <i>nn</i>

Product	Title	Part Number
Expansion cabinet	<i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i>	805-3067- <i>nn</i>
Storage Service Processor	<i>Sun Fire V210 and V240 Server Administration Guide</i>	816-4826- <i>nn</i>
Solaris Operating System	<i>Solaris Handbook for Sun Peripherals</i>	816-4468- <i>nn</i>

Accessing Sun Documentation

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

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

For Sun StorEdge 6920 system documentation, go to:

http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/Midrange/6920/index.html

Related Third-Party Documentation

For installation instructions and other information about Brocade and McData Fibre Channel switches, refer to the product documentation provided by the vendor.

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Contacting Sun Technical Support

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

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

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 6920 System Administration Guide, part number 819-0123-10.

Product Overview

This chapter describes overview information about the Sun StorEdge 6920 system. It contains the following sections:

- [“System Overview” on page 2](#)
- [“Configuration Overview” on page 19](#)
- [“Getting Started With Configuration” on page 26](#)
- [“Monitoring Overview” on page 28](#)
- [“Getting Started With Monitoring” on page 32](#)

System Overview

This section describes the features and components of the Sun StorEdge 6920 system. It contains the following subsections:

- [“About the Sun StorEdge 6920 System” on page 2](#)
- [“About High-Availability Features” on page 4](#)
- [“About the Sun StorEdge 6920 Hardware” on page 4](#)
- [“About the Sun StorEdge 6920 Software” on page 14](#)
- [“About Data Services” on page 17](#)
- [“About User Roles” on page 18](#)

About the Sun StorEdge 6920 System

The Sun StorEdge 6920 system has a modular architecture with integrated system-wide manageability. The Sun StorEdge 6920 system features include:

- Storage virtualization and pooling to simplify storage management.
- Scaling and aggregation of both capacity and performance. System capacity scales from 504 Gbytes to 65 Tbytes and capacity can be added while the system is online.
- Centralized management and monitoring through a browser interface and remote scripting client.
- High availability and failover capability through redundant hardware components.
- Serviceability of components that can be replaced while the system is online.
- Diagnostic monitoring to enhance reliability, availability, and serviceability (RAS) of the system through the Sun Storage Automated Diagnostic Environment.
- Remote monitoring, troubleshooting, and servicing by Sun-trained personnel through Sun StorEdge Remote Response service.
- Virtualization of heterogeneous external storage to facilitate data migration from existing storage devices.

Centralized data services including Sun StorEdge Storage Pool Manager, Sun StorEdge Data Snapshot, Sun StorEdge Data Mirror, and Sun StorEdge Data Replicator. The illustration below depicts the Sun StorEdge 6920 system environment.




FIGURE 1-1 System Environment

About High-Availability Features

The Sun StorEdge 6920 system is designed to provide high levels of reliability, availability, and serviceability (RAS). All components in the system's data path are redundant so that the data path has no single point of failure. If one component fails, the system automatically switches to the alternate component. The following components have been designed for high availability:

- **Management interface cards (MICs)** – Two MICs share one IP address, operating in master/alternate master mode.
- **Storage trays** – The controller trays operate in master/alternate master mode in each array. In addition, the loop card (LPC) and cables each provide one path to devices in a storage tray. If the loop card fails, the data can be accessed through the cables.
- **Hot-spare drives** – Hot-spare drives are always powered up and available but are not part of the array's virtual disk. Two types of hot-spare drives are available:
 - **Array hot-spare** – One or more disks are configured as part of a storage pool. The pool's array hot-spare is available to all virtual disks in the array.
 - **Dedicated hot-spare** – A disk serves as a hot-spare for a designated virtual disk.

A configuration that has dedicated hot-spares in every storage tray and array hot-spares in every pool supports high availability.

The system uses several mechanisms to recognize a hardware failure. When a component recognizes that it has failed or the management software detects a failure, the component is removed from active service and its functions are assigned to other components of the system. The system generates an error event and records it in log files located on the Storage Service Processor.

About the Sun StorEdge 6920 Hardware

This section describes system hardware components. It contains the following subsections:

- [“About the System Cabinet” on page 5](#)
- [“About the Service Processor Panel” on page 6](#)
- [“About the Storage Service Processor” on page 8](#)
- [“About the Storage Service Processor Accessory Tray” on page 9](#)

- “About the Data Services Platform” on page 9
- “About Storage Arrays” on page 12

About the System Cabinet

The system base cabinet contains all the hardware components for the system, as shown in the following figure.




FIGURE 1-2 System Cabinet

From the front of the system base cabinet, you can view the LEDs for the Storage Service Processor, Data Services Platform (DSP), and storage arrays. Each of the 14 Fibre Channel RAID disk drives has LEDs that indicate drive activity:

- If the LEDs are green, the drives are ready to send and receive input/output (I/O) activity.
- If the LEDs are amber, there is a fault.
- If the LEDs are blue, you can safely remove the drive.

From the back of the base cabinet, you can access the service panel for cabling data hosts, storage area network (SAN), local area network (LAN), and external storage devices to the system. You can also access storage arrays and the DSP's field replaceable units (FRUs) such as power and cooling units, array controllers, and management interface cards (MICs).

Depending on your configuration, your system can consist of only a base cabinet, or it can include one or two expansion cabinets.

About the Service Processor Panel

The Service Processor panel simplifies the cabling to the system. The service panel provides the following connections:

- Modem connection
- LAN connections for management
- Serial ports
- FC I/O connections for data hosts, external storage, and FC-based remote replication
- Gigabit Ethernet connections for Gigabit Ethernet-based remote replication
- Power connections for expansion cabinet management
- Ethernet and Fibre Channel (FC) connections for expansion cabinets

An auxiliary port, AUX, to support an external modem for use if the system is used at a location where the telephone system does not support the internal modem. The following figure shows the service panel.




FIGURE 1-3 Service Panel

You connect cables to these accessible panel connections rather than to individual components of the system. When the system has expansion cabinets, the cabling between the base and expansion cabinets is also handled through the service panels. The following figure shows the cabling between the base cabinet and a single expansion cabinet.

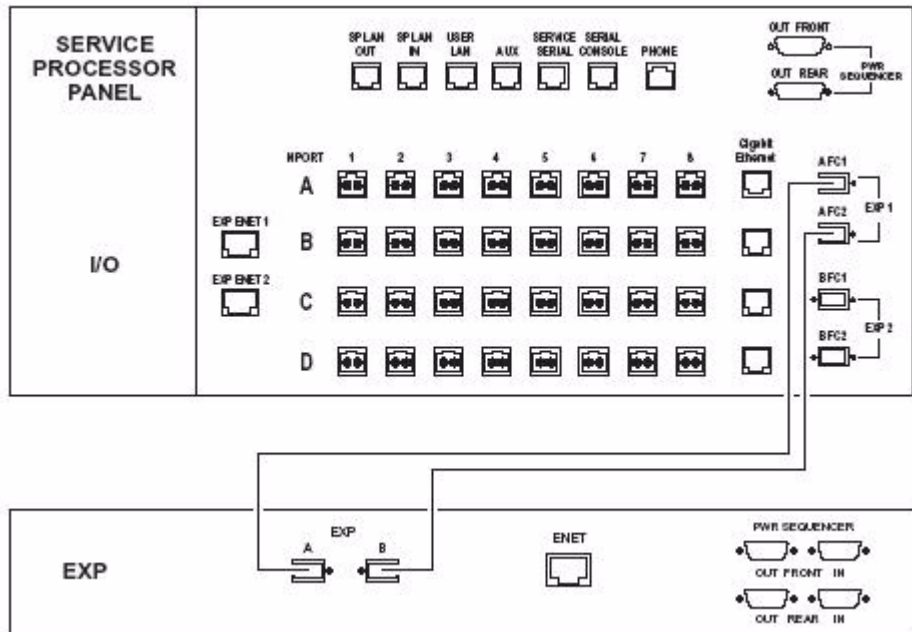


FIGURE 1-4 Single Expansion Cabinet

The following figure shows the cabling between the base cabinet and two expansion cabinets.




FIGURE 1-5 Two Expansion Cabinets

About the Storage Service Processor

The Storage Service Processor is a management host residing in the system's base cabinet. The management software arrives pre-installed on the Storage Service Processor. The Storage Service Processor is cabled to the service panel, allowing for easy serial and local area network (LAN) connections. You do not connect directly to the Storage Service Processor.

User configuration data from the Storage Service Processor is duplicated on the Storage Service Processor's internal hard disk and an external USB flash disk. The USB flash disk is connected to the Storage Service Processor accessory tray, which

stores system information for the Storage Service Processor and Data Services Platform (DSP). You insert the flash disk into the Storage Service Processor through a USB port.

If a Storage Service Processor fails and needs to be replaced, the USB flash disk from the failed Storage Service Processor can be removed and inserted into the replacement Storage Service Processor to restore the configuration settings.

About the Storage Service Processor Accessory Tray

The Storage Service Processor accessory tray provides a number of functions for the system, including:

- **Network terminal concentrator (NTC)** – The NTC provides a modem connection point for the Sun StorEdge Remote Response service.
- **Modem** – The modem connects to the Sun StorEdge Remote Response service and allows the Sun StorEdge 6920 system to communicate with Sun Solution Centers over a dial-up telephone connection.
- **Router with a firewall** – The router provides an additional layer of security for Sun by using network address translation (NAT).
- **Ethernet hub** – The Ethernet hub provides Ethernet connectivity for all components of the Sun StorEdge 6920 system. The Ethernet hub in the base cabinet is located under the Storage Service Processor accessory tray. The Ethernet hub in the expansion cabinet is mounted on the rear door of the cabinet.
- **Flash disk connection** – The USB flash disk connection is used for backing up system configuration files.
- **Integrated power supply module** – There is one power connection to the Ethernet hub, Storage Service Processor, and Storage Service Processor accessory tray.

About the Data Services Platform

The Data Services Platform (DSP) provides virtualization services for the storage arrays in the system.

The DSP's configuration information is stored in a resident flash disk, which can be moved to a new DSP if the current DSP ever needs replacement. The following figures show the front and back views of the DSP.

Front View:




FIGURE 1-6 Front View of the DSP

Back View:




FIGURE 1-7 Back View of the DSP

The DSP includes the following components:

- **Power supplies** – The DSP has two bays located at the top of the rear of the chassis, each containing a 750-watt AC/DC power supply. These power supplies provide a nominal 48V DC of power distributed across the chassis midplane. They are hot-swappable and provide redundancy. It takes one power supply to power a fully loaded chassis.

The power supplies have dedicated cooling fans, which draw ambient air in from the front of the unit and out at the back. Each unit includes an integral thermal shutdown circuit in case of fan failure.

- **Power cord** – A detachable power cord is connected to each power supply. Latching handles prevent removal from the chassis until power is disconnected.
- **Fan assembly** – The fan assembly consists of two variable-speed DC-powered fans and a fan controller card (FCC). The fans draw ambient air into the chassis through the front intake vents and draw direct heated exhaust from the fan assembly. This configuration eliminates recirculation of exhaust air generated by other units in the same rack or adjacent racks. The fan assembly is hot-swappable from the front of the system chassis.
- **Management interface cards (MICs)** – MICs are responsible for providing the system interface and management control functions for the DSP chassis. Each MIC has the following:
 - High-performance microprocessor
 - RJ-45 serial port
 - 10/100 Ethernet port
 - PCMCIA device (a PC card) for storing operational images and configuration storage
- **Storage resource card (SRC) sets** – The SRC set handles packet processing for the DSP's FC interfaces. Each card set consists of two boards that are interconnected through the chassis midplane:
 - The SRC, the line-processing card, plugs into the front of the DSP chassis. Each processor on the SRC manages two corresponding ports on the storage I/O (SIO) card.
 - The SIO card plugs into the rear of the DSP chassis. The SIO card provides the FC ports for the DSP. This card is always paired with an SRC. The Sun StorEdge 6920 system supports two types of SIO cards. The SIO-8 card has eight FC ports and the SIO COMBO card has six FC ports and one gigabit Ethernet port (pictured below).




FIGURE 1-8 SIO COMBO Card

- **SRC mezzanine** – The SRC mezzanine provides a local memory cache for the processors and is equipped with 2 gigabytes of memory.
- **Switch fabric card (SFC)** – The SFC provides the central data path switching function for the DSP chassis.

About Storage Arrays

Storage arrays supply the physical storage of the system. The storage array design is modular, with a variety of possible configurations. Each configuration option includes two controllers (also referred to as a controller pair) to provide redundancy and failover capabilities. Each storage array also has redundant Fibre Channel (FC) data paths and two power supplies with an integral battery backup system. In the event of a total power failure, each array has sufficient power from the batteries to shut down the system in an orderly fashion. In addition, each disk drive has system area information that is mirrored across all drives for redundancy so that data can be recovered from any functional drive.

Each storage tray can contain 7 or 14 disk drives. A storage tray with fewer than 14 disk drives can be upgraded.

The system is available in three supported storage array configurations.

TABLE 1-1 Supported Configurations

Option	Controllers x Trays	Number of Disk Drives	Minimum Capacity	Maximum Capacity
1	2 x 2	14 to 28	504 gigabytes	4 terabytes
2	2 x 3	28 to 56	1 terabyte	8 terabytes
3	2 x 4	42 to 84	1.5 terabytes	12 terabytes

The following figure shows each of these configurations.

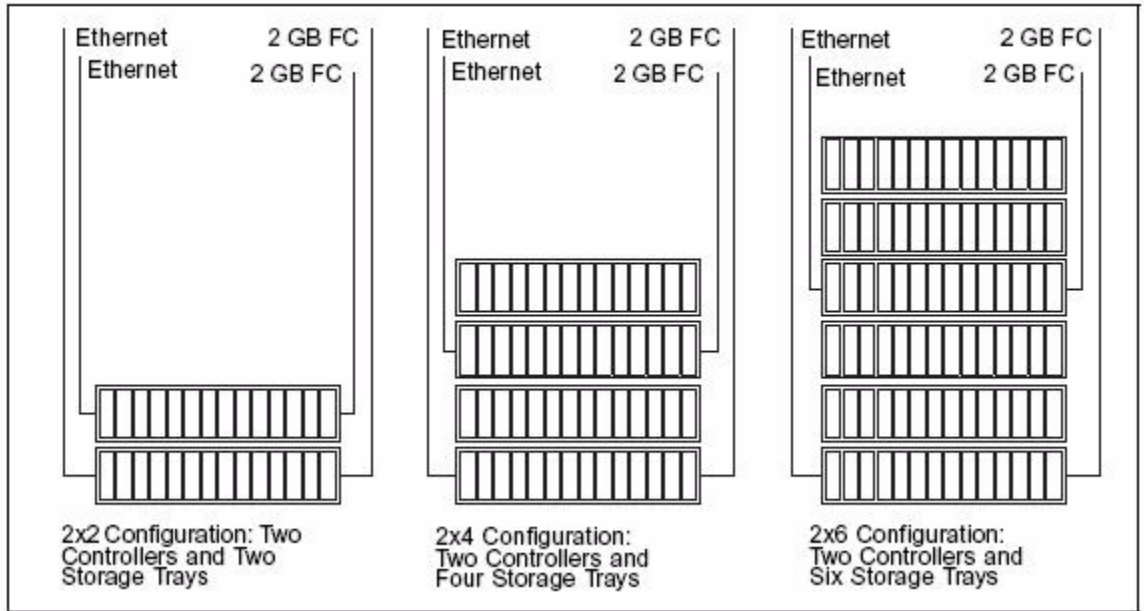


FIGURE 1-9 Supported Configurations

The configuration of the system is based on requirements for throughput, amount of storage, and economy. For example, if throughput is more important to you than economy, your system might use the option 1 storage array configuration. In contrast, if the amount of storage is more important than throughput, your system might use the option 2 or 3 storage array configuration.

A controller tray contains disk drives and built-in RAID management hardware. An expansion tray contains disk drives only and is managed by a controller tray.

FC I/O connections on the service panel connect to the controller of each array in the base cabinet. This cabling is attached at the factory before the system is shipped.

In addition to having internal storage arrays, the Sun StorEdge 6920 system provides support for virtualization of data stored on external storage devices. Data stored on external storage devices can be preserved and added to the system as a legacy volume. You can also use external storage as a means of increasing storage capacity. In this case, the data on the external storage device is not preserved and the capacity is added to a storage pool as one volume. You can also migrate data from an external storage device to a volume on the system using data mirroring.

About the Sun StorEdge 6920 Software

This section describes the system's main software components. It contains the following subsections:

- [“About User Interfaces” on page 14](#)
- [“About Resident Software” on page 14](#)
- [“About Required Host Software” on page 15](#)

About User Interfaces

You can access the system using the two main interfaces to the system:

- A browser interface, the primary interface for configuring, managing, and monitoring the system
- A command-line interface (CLI), available through a remote scripting CLI client for remote management

The remote scripting CLI client provides the same control and monitoring capability as the browser interface, and it is scriptable for running frequently performed tasks. If you want to configure storage from a remote management host, you can install the remote scripting CLI client on an external server and then use the CLI commands to access the system. You can access the remote scripting CLI client from any management host that is connected to the user local area network.

The remote scripting CLI client for the Solaris Operating System is provided on the Host Installation Software CD that is delivered with the system. Host software for other operating systems is available from the Sun Download Center (SDLC).

About Resident Software

The following software is pre-installed on the Sun StorEdge 6920 system:

- **System and storage management software** – Provides a centralized, web-based management platform, that provides configuration and management services through a centralized management tool.
- **System monitoring and diagnostic software** – Provides a diagnostic monitoring tool for the system. It can be configured to monitor on a 24-hour basis, collecting information that enhances the reliability, availability, and serviceability (RAS) of the Sun StorEdge 6920 system.

- **Revision maintenance software** – Provides a tool for upgrading firmware revisions of all components resident in the system. It also keeps track of component firmware levels and firmware revisions.
- **Sun StorEdge Remote Response service software** – Connects to the Sun Service Center, allowing trained personnel to remotely monitor, troubleshoot, diagnose, and service the system. If the Sun Service Center organization determines that there is a problem, Sun will recommend and perform maintenance on the system.
- **Data services software** – Facilitates the effective management of data and provides critical data services for key applications in the data center. The data services applications are available as optional licensed software.

About Required Host Software

In addition to the software that is pre-installed on the system, there is additional software that must be installed on hosts for the system either to function properly or to gain additional functionality.

Solaris Host Software

The Host Installation Software CD that comes with the system includes a variety of Solaris host-based software for managing the Sun StorEdge 6920 system. You specify the functions that you require, and the CD installs the corresponding software.

The software on the Host Installation Software CD includes the following:

- **Sun StorEdge SAN Foundation software** – Incorporates drivers and utilities that enable Solaris data hosts to connect to, monitor, and transfer data in a storage area network.

To enable multipathing functionality, each Solaris data host must have Sun SAN Foundation software installed before it can reliably communicate with the system's storage. You may alternatively choose to purchase VERITAS software to enable Dynamic Multipathing (DMP).
- **Storage Automated Diagnostic Environment Enterprise Edition software** – Resides on an external management station to monitor devices in the SAN.
- **Sun StorEdge Remote Scripting CLI** – Enables remote storage configuration from Solaris hosts for the Sun StorEdge 6920 system.

Host Software for Other Operating Systems

Additional host software enables data hosts that run operating systems other than Solaris to communicate with the Sun StorEdge 6920 system. For details about the supported operating systems and versions, see the *Sun StorEdge 6920 System Release Notes*.

This software consists of the following:

- **Sun StorEdge Traffic Manager Software** – Is required for hosts to make use of the Sun StorEdge 6920 system's storage. This multipathing software incorporates kernel drivers and utilities that enable data hosts to connect to, monitor, and transfer data in a SAN. The Sun StorEdge Traffic Manager software is available for Microsoft Windows 2000, Microsoft Windows 2003, Red Hat Linux, HP-UX, and IBM AIX operating environments.
- **Sun StorEdge Remote Scripting CLI** – Enables remote storage configuration from hosts for the Sun StorEdge 6920 system. The remote scripting CLI client is available for Microsoft Windows 2000, Microsoft Windows 2003, Red Hat Linux, HP-UX, and IBM AIX operating environments.

You can download host software for operating systems other than Solaris from the Sun Download Center at:

http://www.sun.com/software/download/sys_admin.html

Go to the Download Center, select the Sun StorEdge 6920 system related software link, and follow the instructions to register and download the host software.

You must zone like operating systems together. For example, if you use Linux and run Linux hosts that are communicating with the Sun StorEdge 6920 system, you must zone the Linux hosts separately from the hosts that run on other operating systems.

Additional Supported Software and Third-Party Software

For information about additional host-based software and third-party software supported by the Sun StorEdge 6920 system, see the *Sun StorEdge 6920 System Getting Started Guide*. For a list of supported releases and versions, see the *Sun StorEdge 6920 System Release Notes*.

About Data Services

Sun StorEdge 6920 data services software facilitates the management of data and provides critical data services for key applications in the data center. The following data services applications are available as optional licensed software:

- **Sun StorEdge Storage Pool Manager software** – Enables virtualization and pooling of storage assets across applications. Licensing for up to 2 terabytes of capacity is provided with the system. Additional right-to-use (RTU) licensing is required for use of larger capacities.
- **Sun StorEdge Data Snapshot software** – Enables you to create point-in-time snapshot copies of a volume. The snapshot copies can be read/write-mounted by applications to secondary storage and used for backup, application testing, or data mining without the need for the primary storage to be taken offline. The right-to-use (RTU) license is issued per system and is based on the capacity of primary storage that will be copied. For more information, see “About Snapshots” on page 103.
- **Sun StorEdge Data Mirror software** – Enables you to create local, independent read-and-write copies of data. Applications can access the mirrored data while other operations, such as backups, data recovery, and application testing occur in parallel without disrupting production application access. Licensing is based on the total capacity of the primary volumes that are being mirrored. For more information, see “About Mirroring” on page 93.
- **Sun StorEdge Data Replicator software** – Enables you to replicate data from one site to another. Replica sites can be located anywhere in the world, and data can be transparently written to both primary and secondary sites simultaneously. The software includes a fast resynchronization feature that allows data to be synchronized quickly in the event of a link failure between sites, or at predetermined intervals. Licensing is based on the total capacity of the primary volumes that are being replicated. For more information, see “About Data Replication” on page 116.

About User Roles

The type of user role you assign to users determines their level of system access. There are three types of user roles, each with a default user name and password.

TABLE 1-2 User Roles and Default Names and Passwords

User Role	Default User Name	Default Password	Description
Administrative	admin	!admin	The administrative role has read/write privileges for system-wide administrative settings. For example, a user assigned the administrative role can modify users, licenses, and system attributes.
Storage	storage	!storage	The storage role has all administrative privileges as well as read/write privileges for all storage configuration settings. For example, a storage user has all of the privileges of the administrative role as well as full access to array configuration and monitoring.
Guest	guest	!guest	The guest role has read-only privileges and can only view certain information. A guest user cannot modify any settings or features.

Multiple users assigned to the administrative or storage roles can be logged in concurrently. However, because storage and administrative users have write privileges, there is a risk of one user's changes overwriting another user's previous changes. Therefore, you should develop policies about who can make changes and how to notify others.

Configuration Overview

This section describes the default configuration of the system and factors to consider before you change the default configuration. It contains the following subsections:

- [“About the Default State of the System” on page 19](#)
- [“About Provisioning Storage” on page 21](#)
- [“About Physical Storage Elements” on page 24](#)
- [“About Logical Storage Elements” on page 25](#)

About the Default State of the System

After you complete the basic configuration tasks described in the *Sun StorEdge 6920 System Getting Started Guide*, your system has the following characteristics:

- The Ethernet connection is cabled to a host or a gateway on the local area network (LAN).
- A service console is connected.
- The system is powered on.
- The IP address and netmask have been configured.
- Initial setup of the system has been completed.
- One data host, an initiator, is assigned to the DEFAULT storage domain.
- The DEFAULT storage domain uses one storage pool named Default.
- All storage ports are assigned to the Default storage pool.

If you are upgrading an existing system, the configuration is dependent on the system’s prior configuration history. The Default storage profile, which the Default storage pool uses, has the specifications described in the following table.

TABLE 1-3 Default Storage Profile Specifications

Storage Attributes	Value
RAID level	RAID-5
Segment size	64 KB
Dedicated hot-spare	Not Enabled
Read-ahead mode	Enabled
Number of drives	Variable

TABLE 1-3 Default Storage Profile Specifications (Continued)

Storage Attributes	Value
Array type	Best Available Match - Bandwidth
Virtualization strategy	Stripe
Stripe size	1 MB

Your system ships from the factory with the default configuration described in the following table.

TABLE 1-4 Default Storage Configuration

Default Configuration		Configuration Options
Storage domains	<p>All storage elements are in one storage domain (DEFAULT), which can be accessed by all hosts (initiators).</p> <p>This storage domain contains the Default storage pool with its associated disks and virtual disks.</p>	<p>You can choose either of the following options:</p> <ul style="list-style-type: none"> • Use the existing default configuration (as is) • Create a custom storage domain
Storage profiles	<p>The default storage domain uses the Default profile to specify its storage parameters.</p>	<p>You can choose any of the following options:</p> <ul style="list-style-type: none"> • Use the existing default configuration (as is) • Select one of the predefined storage profiles shipped with the system • Create a custom storage profile <p>Note: You cannot delete or modify any of the factory-provided profiles, including the Default profile, nor can you delete or modify a profile that is in use and associated with a storage pool.</p>
Storage pools	<p>All virtual disks are in one storage pool (Default).</p>	<p>You can choose either of the following options:</p> <ul style="list-style-type: none"> • Use the existing default configuration (as is) • Create a new storage pool
Virtual disks	<p>Virtual disks are preconfigured at the factory, based on the number of drives ordered for your system. The virtual disks are configured as RAID-5. If an array has 7 drives per tray, all 7 drives are configured as a single RAID-5 virtual disk. If an array has 14 drives, the tray is configured as two RAID-5 virtual disks.</p>	<p>You can choose any of the following options:</p> <ul style="list-style-type: none"> • Use the existing default configuration (as is) • Move an existing, unused virtual disk to a different storage pool with similar attributes. • Delete the existing RAID-5 virtual disk and reconfigure the recovered disks into a new virtual disk in a storage pool with a different storage profile.

About Provisioning Storage

In the default storage configuration, all data hosts are in one storage domain and all available storage is in one storage pool. Any host (initiator) in the domain can access any storage in the pool. However, your organization's needs determine how you adjust the storage configuration. Instead of assigning hosts to specific physical storage, you provision the storage for your organization by creating domains of initiators and pools of virtual storage.

In addition to domains and pools, the system has physical and logical storage elements that you use to provision your storage:

- **Physical storage elements** - Arrays internal to the base system or expansion cabinet, arrays external to the base system or expansion cabinet, storage trays, and disks
- **Logical storage elements** - Volumes, replication sets, virtual disks, pools, profiles, snapshots, domains, and external storage volumes

The relationships among the physical and logical storage elements is shown in the following diagram.

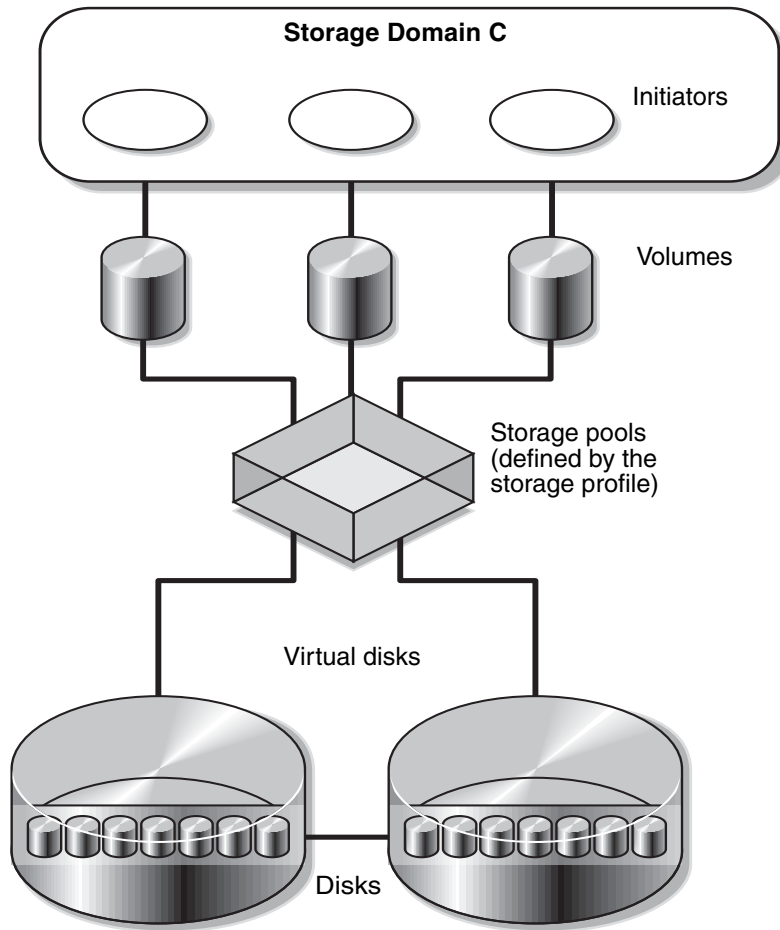


FIGURE 1-10 Relationship Among Physical and Logical Storage Elements

Before you can allocate storage appropriately, consider the following requirements for your site:

- **Security** – By creating additional storage domains, you segregate initiators (as shown in the following figure). For example, the hosts that handle financial data store their data in a different domain from the domain used by hosts that handle research data.

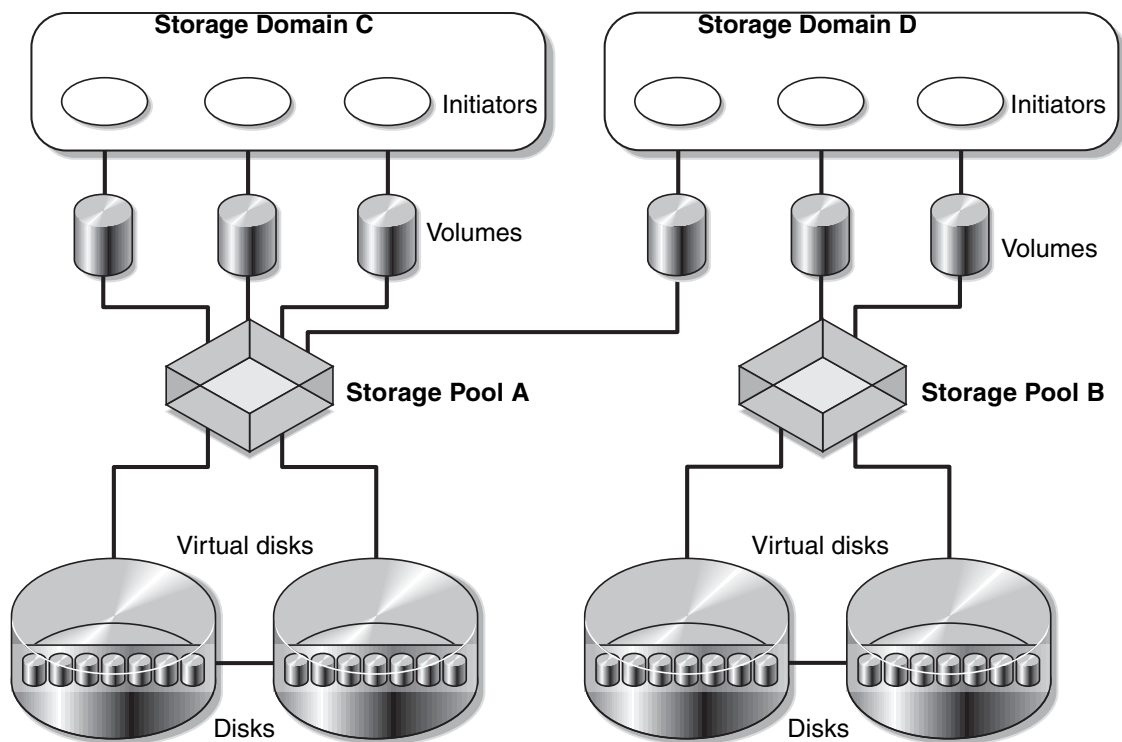


FIGURE 1-11 Example Configuration

- **I/O** – The Default storage profile specifies a general, balanced access to storage, but some parts of your organization might require one or more of the characteristics to be optimized at the expense of other attributes. The system provides a set of profiles to meet various needs. You can also create custom profiles.
- **Performance** – In general, the number of host ports is equal to the number of storage ports. You can add more hosts, but performance declines if the number of ports remains the same. To maintain performance, you can expand the number of ports by adding storage resource card (SRC) sets. For information about how to add an SRC set, go to the Sun Java™ Web Console page and click Sun Storage Automated Diagnostic Environment > Service Advisor > X-Options.

To increase performance beyond what a balanced configuration provides, you can do the following:

- Configure the system to include more host ports if the system supports an on-line transaction processing (OLTP) application.

Note: Configuring too many host ports might cause performance to degrade.

- Implement a stripe virtualization strategy so that volumes can be spread over multiple controllers, thus increasing the number of access points and amount of cache available for a specific volume.

Note: Some of these attributes are part of the storage profiles that ship with the system. You can create a custom profile to best fit your environment.

- **Access to external storage devices** – Use an external storage device’s raw storage or user data. The following figure shows that an external storage device can be associated with one domain or multiple domains.

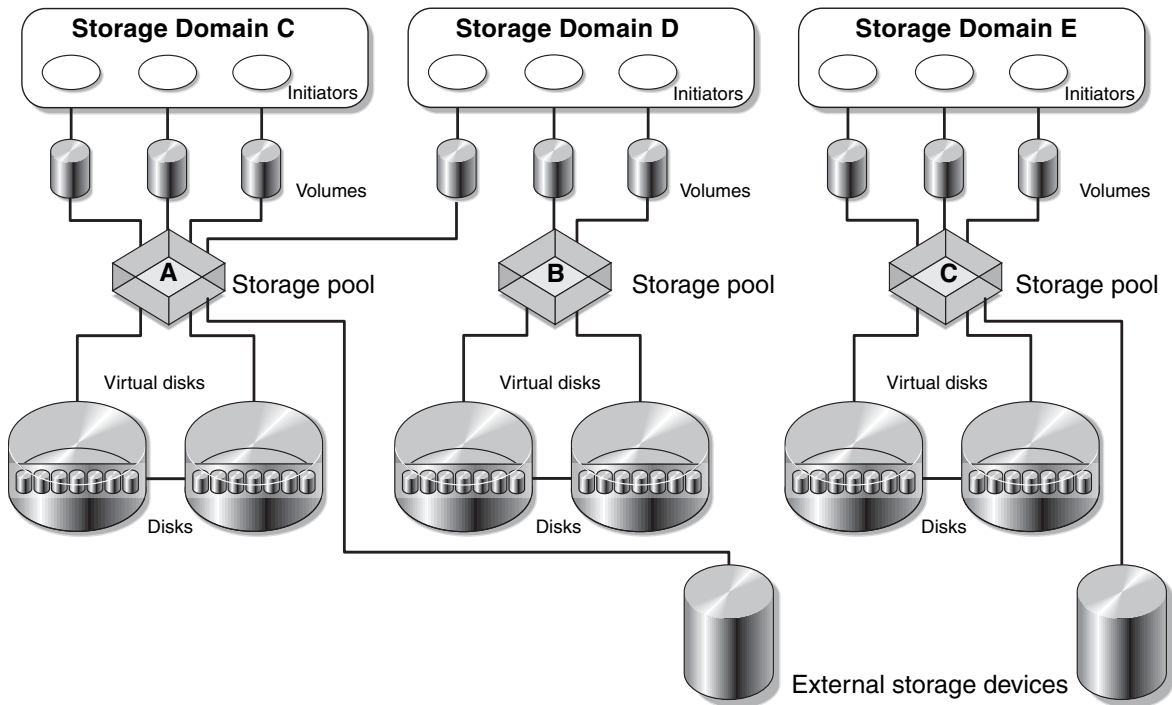


FIGURE 1-12 External Device Being Associated With One or Multiple Domains

About Physical Storage Elements

Consider the following physical storage elements before you decide how to distribute data across the available physical storage:

- Initiators are the protocol-specific physical ports that initiate the input and output (I/O) exchanges with the system.

- Ports connect the system to data hosts, external storage devices, and other Sun StorEdge 6920 systems. The system also has inbound ports that connect the Data Service Platform (DSP) to the arrays in the system or an expansion cabinet.
- In a SCSI environment, an initiator is a SCSI node that sends requests to a SCSI device. SCSI adapters in hosts are usually initiators, but storage devices that provide virtualization can also act as initiators to other SCSI devices.
- Arrays are sets of storage trays.
- Storage trays hold the disk drives, and support their operation. Each storage tray holds from 7 to 14 disk drives.
- Disk drives are non-volatile, randomly addressable, rewriteable data storage devices.
- External storage is a combination of storage and user data that does not reside within the base cabinet or an expansion cabinet; however, the storage and user data are available to the system from an external storage device. You can use the external storage device for added capacity, or you can migrate the user data to internal storage on the system.

About Logical Storage Elements

Consider the following logical storage elements before you decide how to distribute data across the available physical storage and map it to data hosts:

- Storage pools are collections of similar virtual disks. Each pool has a profile that specifies a set of attributes.
- Virtual disks, also called RAID sets, are not physical disk drives. A virtual disk is a collection of one or more physical disk drives. It is a member of a storage pool and exhibits the properties defined by the pool's profile. The storage array handles a virtual disk as if it were an actual disk drive.
- Volumes are divisions of a pool, consisting of pairs of virtual disks, and are assigned to initiators in the storage domain.
- Volume snapshots are copies of the data in a volume at a specific moment. Snapshot copies can be made without interrupting the normal operation of the system.
- Replication sets are pairs of volumes that are in physically separate locations and share identical copies of data. You can use replication sets to implement disaster recovery and business continuance strategies.
- A mirror is a volume composed of multiple identical copies of user data, which are known as mirror components. Mirror components are similar to volumes, in that they each reside entirely within a storage pool and are composed of virtual disks. Different components of a mirror can reside in different storage pools.

Getting Started With Configuration

Your system is shipped with a default configuration that simplifies storage provisioning. See “About the Default State of the System” on page 19 for more details.

If you decide to use the default configuration that shipped with the system, you must complete the following tasks:

1. Create volumes within the Default storage pool, as described in “Creating a Volume” on page 89.
2. Assign each initiator in your environment to one of the volumes, as described in “Mapping a Volume to Initiators” on page 90.

If you decide the default configuration does not meet your storage needs, you can change it by doing any of the following:

- Create one or more storage domains

For example, you might want to create a separate domain to isolate financial data and ensure that only specified hosts can access it. See “Configuring Storage Domains” on page 167 for information about using, planning for, and creating new storage domains.

- Create one or more storage pools

For example, if you plan to create a legacy volume from external storage, you might want to create a custom legacy pool using the Legacy profile, rather than use the default legacy pool provided by the system. See “Configuring Storage Pools” on page 153 for information about using, planning for, and creating storage pools and “Configuring Storage Profiles” on page 158 for information about using, planning for, and creating storage profiles.

- Move a virtual disk to another pool, or delete a virtual disk and create a new virtual disk with another configuration

For example, although the virtual disks are configured as RAID-5 by default, you might want to use a profile with a RAID level of RAID-1. You would need to delete the virtual disk and create a new storage pool with a different profile. See “Configuring Virtual Disks” on page 150 for information about managing virtual disks.

- Access storage and user data that resides on an external storage array or system

For example, you might want to migrate user data from an external storage device to internal storage on the system. See “Managing External Storage” on page 191 for information about using, planning for, and accessing external storage.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Monitoring Overview

This section describes the Sun Storage Automated Diagnostic Environment. It contains the following subsections:

- [“How the Monitoring Software Works” on page 28](#)
 - [“About Monitoring Strategy” on page 29](#)
 - [“About the Monitoring Cycle” on page 29](#)
 - [“About the Event Life-Cycle” on page 30](#)
-

How the Monitoring Software Works

The monitoring and diagnostic software is a distributed application used to monitor and diagnose Sun storage products, Sun-supported switches, and Sun virtualization products. The primary functions of the software are:

- Device health monitoring
- Event and alarm generation
- Notification to configured recipients
- Topology presentation
- Diagnostics
- Revision checking
- Device and device component reporting

The software depends on an agent installed on the Storage Service Processor. The monitoring and diagnostics server package adds a cron entry to the Storage Service Processor. This entry is used to provide periodic probing of the devices. An Hypertext Transfer Protocol (HTTP) or Secure HTTP application service is also added during installation.

The agent is at configured intervals to probe devices and monitor log files. A configuration file maintains the inventory of devices that the installed agent monitors. Events are generated with content, such as probable cause and recommended action, to help facilitate isolation to a single field-replaceable unit (FRU).

About Monitoring Strategy

The following procedure is a typical strategy for monitoring.

1. Monitor the devices.

To get a broad view of the problem, the site administrator or Sun personnel can review reported information in context. This can be done by:

- Displaying the device itself
- Displaying the topology
- Analyzing the device's event log

2. Isolate the problem.

For many alarms, information regarding the probable cause and recommended action can be accessed from the alarm view. In most cases, this information enables you to isolate the source of the problem. In cases where the problem is still undetermined, diagnostic tests are necessary.

From the browser interface, you can execute tests remotely using the agent. Once the problem is fixed, you can clear the alarms for the device.

About the Monitoring Cycle

Agent execution is controlled by the `cron` daemon on each server. The high-level steps of a monitoring cycle are as follows.

1. Verify that the agent is idle.

The system generates instrumentation reports by probing the device for all relevant information, and it saves this information. The system then compares the report data to previous reports and evaluates the differences to determine whether health-related events need to be generated.

Events are also created from information found in log files. For example, all errors and warnings are translated into a log event without further analysis. Most events are generated because a rule or policy in the software concluded that a problem exists, but if the storage array indicates issues in the `syslog` file, an event is immediately generated.

Note: Aggregated events and events that require action by service personnel (known as actionable events) are also referred to as alarms.

2. Store instrumentation reports for future comparison.

Event logs are accessible from the Administration tab of the user interface. The software updates the database with the necessary statistics. Some events require that a certain threshold be attained before an event is generated. For example, having the cyclic redundancy count (CRC) of a switch port increase by one is not sufficient to trigger an event, since a certain threshold is required.

3. Send the events or alarms to interested parties.

Events are sent only to recipients that have been set up for notification. The types of events can be filtered so that only pertinent events are sent to each individual.

Note: If they are enabled, the email providers and the Sun Network Storage Command Center (NSCC) receive notification of all events.

About the Event Life-Cycle

Most storage network events are based on health transitions. For example, a health transition occurs when the state of a device goes from online to offline. It is the transition from online to offline that generates an event, not the actual offline value. If the state alone were used to generate events, the same events would be generated repeatedly. Transitions cannot be used for monitoring log files, so log events can be repetitive. To minimize this problem, attach thresholds to entries in the log files.

The software includes an event maximums database that keeps track of the number of events generated about the same subject in a single eight-hour time frame. This database prevents the generation of repetitive events. For example, if the port of a switch toggles between offline and online every few minutes, the event maximums database ensures that this toggling is reported only once every eight hours instead of every five minutes.

Event generation usually follows this process:

1. The first time a device is monitored, a discovery event is generated. It is not actionable but is used to set a monitoring baseline, primarily for the Network Storage Command Center (NSCC). When NSCC is enabled, a discovery event is generated to re-establish the monitoring baseline. This event describes, in detail, the components of the storage device. Every week after a device is discovered, an audit event is generated with the same content as the discovery event.
2. A log event can be generated when interesting information is found in host or storage log files. This information is usually associated with storage devices and sent to all users.

3. Events are generated when the software detects a change in the content of the instrumentation report, probes the device, and compares the report to the last instrumentation report, which is usually only minutes old. `StateChangeEvent` and `ValueChangeEvent` categories represent most of the events that are generated.
4. When possible, the agent combines events to generate aggregated events.

Note: Aggregated events and events that require action by service personnel (known as actionable events) are also referred to as alarms. Some alarms are based on a single state change and others are a summary of events where the event determined to be the root cause is advanced to the head of the queue as an alarm. The supporting events are grouped under the alarm and are referred to as aggregated events.

Getting Started With Monitoring

The following table lists the tasks and operations required to begin using the product. For an overview of the software, see “Monitoring Overview” on page 28.

TABLE 1-5 Monitoring Tasks

Task	For More Information
<i>Required Tasks</i>	
Verify site information. Ensure that company and contact information is correct.	“Displaying and Editing Site Setup Information” on page 67
Set up email notification. Specify email addresses for local email and pager notification recipients.	“Configuring Email Notification” on page 72 “Adding an Email Notification Recipient” on page 73
Set up remote notification. Enable and configure remote Sun notification recipients, including Network Storage Command Center (NSCC), Sun StorEdge Remote Response (SSRR), and Simple Network Management Protocol (SNMP) traps.	“Configuring Remote Notification” on page 76
Run agent. Manually initiate an agent run to verify system health.	“Running the Agent” on page 79
Review system topology. Display and review system topology to verify expected configuration.	“Displaying the System Topology” on page 225

Using the Browser Interface

This chapter describes how to use the browser interface. It contains the following sections:

- [“About the Browser Interface” on page 34](#)
- [“About the Page Layout” on page 34](#)
- [“Controlling the Display of Table Information” on page 38](#)
- [“Using Forms” on page 41](#)
- [“Searching for System Elements” on page 41](#)
- [“Using Help” on page 42](#)
- [“Logging Out of the Management Software” on page 43](#)

About the Browser Interface

The browser interface provides you with an easy-to-use interface to configure, manage, and monitor the system. You navigate through the browser interface as you would a typical web page. You use the navigational tabs to move among pages within an application. You can click a link to get details about a selected item. You can also sort and filter information displayed on a page. When you place your pointer over a button, tab, link, icon, or column, a tooltip provides a brief description of the object.

Each page uses a form or table format to display data.

About the Page Layout

Each page of the browser interface is divided into the following three sections:

- [“About the Page Banner” on page 34](#)
- [“About the Navigational Tabs” on page 36](#)
- [“About the Page Content Area” on page 38](#)

About the Page Banner

Across the top of each page, the banner displays buttons, links, system information, alarm status, and the name of the application.

The following table describes the contents of the banner. Note that the information varies slightly, depending on which application you open.

TABLE 2-1 Contents of the Banner








Button	Description
	Returns you to the Java Web Console page, where you can navigate between the configuration software and the diagnostic software.
	Displays the software version and copyright information.
	<i>Sun StorEdge 6920 Configuration Service only.</i> Refreshes the current page.

TABLE 2-1 Contents of the Banner (*Continued*)

Button	Description
SEARCH	<i>Sun StorEdge 6920 Configuration Service only.</i> Enables you to quickly locate logical and physical elements defined in the system. You select a component and enter a name or World Wide Name (WWN) for the component you want to locate. An asterisk (*) searches for all instances of the selected component. For example, you can search for all initiators or only those initiators that match a specified name or WWN. For more information, see “Searching for System Elements” on page 41.
SERVICE ADVISOR	<i>Sun Storage Automated Diagnostic Environment only.</i> Provides information and procedures to enable you to service hardware components.
PREFERENCES	<i>Sun Storage Automated Diagnostic Environment only.</i> Enables you to set options for the default behavior of the user interface. For more information, see “Setting Diagnostic Viewing Preferences” on page 206.
LOG OUT	Logs you out of the Java Web Console and the current application.
HELP	Opens the onlinehelp in a separate window.
System Information and Status	
User: storage	Displays the name of the user who is currently logged in to the system.
Server: sp1	Displays the name of the system.
Jobs Running: 0	<i>Sun StorEdge 6920 Configuration Service only.</i> Displays the number of configuration jobs currently running. To get more information about the jobs, click the Jobs tab.

TABLE 2-1 Contents of the Banner (*Continued*)

Button	Description
Current Logins: 1	<i>Sun StorEdge 6920 Configuration Service only.</i> Displays the number of users currently logged in to the system. Click the link to open the Active User Summary, which displays the user name, role, client type, and IP address for each logged-in user.
Last Update: Feb 2	Displays the latest date and time that data was retrieved from the server that you are administering. The latest data is collected and displayed each time you refresh the browser window or perform an action in the browser.
Current Alarms:	<p>Displays the current number of each type of alarm. There are four alarm types:</p> <p> Down,  Critical,  Major, and  Minor.</p> <p>To get more information about the alarms, click the Current Alarms link.</p> <p>When you click Current Alarms from the diagnostic environment, the Alarm Summary page is displayed.</p> <p>When you click Current Alarms from a configuration service window, the diagnostic environment is launched in a separate window.</p>

About the Navigational Tabs

You use the navigational tabs to move among pages within an application. The tabs vary, depending on the application that is open.

Sun File System Manager Tabs

The following table describes the Sun Storage Automated Diagnostic Environment tabs.

TABLE 2-2 Sun Storage Automated Diagnostic Environment Tabs

Tab	Description
Alarms	Provides access to alarm summary and alarm detail information.
Inventory	Provides functions for the discovery of new devices, and provides access to summary and detail information for system devices.
Topology	Displays current or archived graphical views of the storage network topology.
Jobs	Displays information on current, completed, and archived monitoring jobs.
Administration	Provides functions for the configuration of system functions and administrative components.

Sun StorEdge 6920 Configuration Service Tabs

The following table describes the Sun StorEdge 6920 Configuration Service tabs.

TABLE 2-3 Sun StorEdge 6920 Configuration Service Tabs

Tab	Description
Logical Storage	Enables you to configure volumes, snapshots, replication sets, virtual disks, storage pools, storage profiles, and storage domains.
Physical Storage	Enables you to configure initiators, ports, arrays, trays, disks, and external storage devices.
External Storage	Enables you to view external storage devices that have been discovered by the system and import the devices for use by the system.
Jobs	Provides access to current and historical configuration job information.
Administration	Provides functions for the configuration of system functions and administrative components.

About the Page Content Area

The content section of each page displays storage or system information as a form or table. You click navigational tabs to move among pages.

Controlling the Display of Table Information

Tables display data in a tabular format. You can use the following objects to control the display of data on a page.

TABLE 2-4 Table Objects











Control/Indicator	Description
	<p>Enables you to display only the information that interests you.</p> <p>When filtering tables, follow these guidelines:</p> <ul style="list-style-type: none"> • A filter must have at least one defined criterion. • A filter applies to the current server only. You cannot apply a filter to tables across multiple servers. <p>To filter a table, choose the filter criterion you want from the table's Filter drop-down menu.</p>
	<p><i>Events page only.</i> Enables you to filter event properties that are displayed in the Events table. You can filter by agent host, category, severity, events, or event type. You can also specify whether the display is to include aggregated events and indicates the number of Kbytes that you want read from the events log file.</p>
	<p>Enable you to toggle between displaying all rows and displaying 15 or 25 rows one page at a time. When the top icon is displayed on a table, click the icon to page through all data in the table. When the bottom icon is displayed in a table, click the icon to page through 15 or 25 rows of data.</p>
	<p>Enable you to select or deselect all of the check boxes in the table. Use the icon on the left to select all of the check boxes on the current page. Use the icon on the right to clear all of the check boxes on the current page.</p>






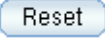
TABLE 2-4 Table Objects (*Continued*)

Control/Indicator	Description
	<p>Indicates that the column in the table is sorted in ascending order. The ascending sort order is by number (0-9), by uppercase letter (A-Z), and then by lowercase letter (a-z).</p> <p>Click this icon to change the sort order of the column to descending.</p> <p>A closed icon indicates the column by which the table is currently sorted.</p>
	<p>Indicates that the column in the table is sorted in descending order. The descending sort order is by lowercase letter (z-a), by uppercase letter (Z-A), and then by number (9-0).</p> <p>Click this icon to change the sort order of the column to ascending.</p> <p>A closed icon indicates the column by which the table is currently sorted.</p>
	<p><i>Alarm Summary page only.</i> Enables you to perform a sort of display information on up to three columns. For example, you can sort the current alarms by date, by severity, and then by type. A closed arrow indicates the column by which the table is sorted.</p>
	<p>Enables you to select the entries that you want to display. Click the button on the left to display the first 25 table entries. Click the button on the right to display the previous 25 table entries.</p>
	<p>Click the button on the left to display the next 15 or 25 table entries. Click the button on the right to display the last 15 or 25 table entries.</p>
	<p>Indicates how many pages are in the table, and displays the page you are currently viewing. To view a different page, type the page number in the Page field and click Go.</p>

Using Forms

Forms have menus, buttons, links, and text fields that allow you to select available options and enter information on a page. The following table describes these elements.

TABLE 2-5 Form Controls

Control/Indicator	Description
	Indicates that you must enter information in this field.
	Lists options from which you can make a selection.
	Displays the part of the form that is indicated by the text next to this icon.
	Returns you to the top of the form.
	Saves the selections and entries that you have made.
	Sets all page elements to the original selections that were displayed when the page was first accessed.

Searching for System Elements

You can easily locate logical and physical elements of the system by using the search feature located in the banner of any configuration service page.

You can search for all elements of a selected type for particular elements that match a specified term. For example, you can search for all initiators or you can search for only the initiators that contain a specific WWN.

To use the search feature:

1. Click Sun StorEdge 6920 Configuration Service.

2. In the banner, click Search.

The Search window is displayed.

3. Select the type of component you want to locate. You can search for arrays, disks, initiators, storage pools, storage profiles, trays, virtual disks, volumes, or all system elements.

4. If you want to narrow your search, enter a term in the text field.

- All elements that contain the specified term in the name or description field will be located. For example, the term “primary” will locate elements with the name of primary, demoprimary, primarydemo, and firstprimarylast.
- The search feature is not case-sensitive. For example, the term “primary” will locate elements that contain primary, Primary, PRIMARY, priMARY, and any other case combination.
- Do not embed spaces or special characters in the search term.
- Use the wildcard (*) only to search for all elements of a selected type. Do not use the wildcard with the search term. If you do, the system will search for the asterisk character.

5. Click Search.

The result of your search is displayed.

6. Click Back to return to the previous page.

Using Help

To view additional information about the configuration or diagnostic software, click Help in the banner of the web browser. The help window consists of a Navigation pane on the left and a Topic pane on the right.




To display a help topic, use the Navigation pane's Contents, Index, and Search tabs. Click the Search tab and click Tips on Searching to learn about the search feature.

TABLE 2-6 Help Tabs

Tab	Description
Contents	Click a folder icon to display subtopics. Click a page icon to display the help page for that topic in the Topic pane.
Index	Click an index entry to display the help page for that topic.
Search	Type the words for which you want to search and click Search. The Navigation pane displays a list of topics that match your search criteria in order of relevancy. Click a topic link to display the help page for that topic. Click the Tips on Searching link for information about how to improve your search results. To search for a particular word or phrase within a topic, click in the Topic pane, press Ctrl+F, type the word or phrase for which you are searching, and click Find.

The Help window icons the have the following meanings.

TABLE 2-7 Help Icons

Control/Indicator	Description
	Click to go back to the previous help topic that you viewed in the current session.
	Click to go forward to the next help topic that you viewed in the current session.
	Click to print the current help topic.

Logging Out of the Management Software

To log out of the software, click Log Out in the window banner.

Performing Connection and Administration Tasks

This chapter describes connection and administration tasks. It contains the following sections:

- [“Connecting the System” on page 46](#)
- [“Specifying General Administrative Settings” on page 63](#)
- [“Performing Monitoring Administration” on page 67](#)

Connecting the System

This section describes setting up network connections and other administrative tasks. It contains the following subsections:

- [“About Network Domains” on page 46](#)
- [“About LAN Names and IP Addresses” on page 47](#)
- [“About Internal System IP Addresses” on page 48](#)
- [“About Network Configuration Options” on page 50](#)
- [“Managing From the Site LAN” on page 51](#)
- [“About Allocating Ports” on page 53](#)
- [“About Host Bus Adapters” on page 55](#)
- [“Connecting an External Switch” on page 56](#)
- [“About Multipathing” on page 57](#)
- [“Administering a System Isolated From the LAN” on page 57](#)
- [“Installing the Remote Scripting CLI Client” on page 58](#)
- [“About Lights-Out Management” on page 58](#)
- [“Powering the System On and Off” on page 59](#)

About Network Domains

The Sun StorEdge 6920 system can be included in three types of network domains:

- Storage Service Processor local area network (LAN), which is used by the Sun StorEdge Remote Response system to remotely monitor the Sun StorEdge 6920 system. When multiple Sun StorEdge 6920 systems (up to eight at the same site) are installed, the systems can share a single telephone line.
- Internal component LAN, in which the Storage Service Processor communicates with components that make up the system. This network is not accessible from outside the system and is isolated from the data path. An Ethernet hub manages internal network traffic on the Sun StorEdge 6920 system.
- Site local area network LAN, which connects to the customer’s internal Ethernet LAN. This network enables access to the system from a management host. A management host is used to configure, control, and monitor the system through a web browser or the remote scripting command-line interface (CLI) client. A firewall is included in the router of the Storage Service Processor accessory tray between the site LAN and the Storage Service Processor LAN.

Note: In addition to the firewall supplied with the Sun StorEdge 6920 system, you can also use a site firewall implemented with your own corporate security policies. However, you cannot use your own firewall in place of the firewall supplied with the Sun StorEdge 6920 system.

About LAN Names and IP Addresses

The following table shows the system's IP addresses and names, which are predefined and reserved.

TABLE 3-1 System's IP Addresses and Names

	Name	IP Address
Storage Service Processor	sp	192.168.0.2
Data Services Platform	dsp00	192.168.0.10

The Storage Service Processor has a second IP address. See "About Internal System IP Addresses" on page 48 for more information.

The IP address and local area network (LAN) name for the Data Services Platform (DSP) resides in the master management interface card (MIC). In the event of a failure, the alternate master MIC assumes the master MIC's IP address and continues its functions.

Each system has at least one array. Up to four arrays can be installed in the base cabinet, and up to six can be installed in an expansion cabinet. Arrays have fixed names and IP addresses based on their locations in the base cabinet and in any expansion cabinets.

The following table lists the reserved IP addresses for all possible array configurations in all possible cabinets.

TABLE 3-2 Reserved IP Addresses For Array Configurations and Cabinets

Description	Name	IP Address
Lowest array (base cabinet 0)	array00	192.168.0.40
Second array from the bottom (base cabinet 0)	array01	192.168.0.41
Third array from the bottom (base cabinet 0)	array02	192.168.0.42
Fourth array from the bottom (base cabinet 0)	array03	192.168.0.43
Lowest array (expansion cabinet 1)	array10	192.168.0.50

TABLE 3-2 Reserved IP Addresses For Array Configurations and Cabinets (*Continued*)

Description	Name	IP Address
Second array from the bottom (expansion cabinet 1)	array11	192.168.0.51
Third array from the bottom (expansion cabinet 1)	array12	192.168.0.52
Fourth array from the bottom (expansion cabinet 1)	array13	192.168.0.53
Fifth array from the bottom (expansion cabinet 1)	array14	192.168.0.54
Sixth array from bottom the bottom (expansion cabinet 1)	array15	192.168.0.55
Lowest array (expansion cabinet 2)	array20	192.168.0.60
Second array from the bottom (expansion cabinet 2)	array21	192.168.0.61
Third array from the bottom (expansion cabinet 2)	array22	192.168.0.62
Fourth array from the bottom (expansion cabinet 2)	array23	192.168.0.63
Fifth array from the bottom (expansion cabinet 2)	array24	192.168.0.64
Sixth array from the bottom (expansion cabinet 2)	array25	192.168.0.65

About Internal System IP Addresses

The Storage Service Processor, internal firewall (also known as a router), and network terminal concentrator (NTC) have their own host names and IP addresses.

The following table lists the names and IP addresses that are reserved for these components and are configured by default.

TABLE 3-3 Default Names and IP Addresses Reserved For These Components

Hardware	IP Address	Host Name
Firewall	10.0.0.242	new_firewall
NTC	10.0.0.243	new_ntc
Storage Service Processor	10.0.0.250	new_sp

If you prefer not to use the reserved names and IP addresses, refer to the following table for the convention to create addresses and host names. If you use the convention, the following table shows the IP addresses and host names of these components established according to this convention for a single system, System 0.

TABLE 3-4 Example Convention for a Single System

Hardware	IP Address n = 2	Host Name N = 0 through 7
Firewall	10.0.0.2	sp0_firewall
NTC	10.0.0.3	sp0_ntc
Storage Service Processor	10.0.0.10	sp0

Note: When you use the Sun StorEdge Remote Response feature, the system that has the telephone line connection must have the sp0 name assigned to it. Other systems can have any unique name from sp1 through sp7.

The naming convention for additional systems increments the IP address by 10, as shown in the following table. In addition to the names and IP addresses for the first system, the table shows the addresses for each additional system, up to the maximum of eight systems on the same internal local area network (LAN).

TABLE 3-5 Naming Convention for Additional Systems

	Firewall Address and Name	NTC Address and Name	Storage Service Processor Address and Name
System 0	10.0.0.2 sp0-firewall	10.0.0.3 sp0-ntc	10.0.0.10 sp0
System 1	10.0.0.12 sp1-firewall	10.0.0.13 sp1-ntc	10.0.0.20 sp1
System 2	10.0.0.22 sp2-firewall	10.0.0.23 sp2-ntc	10.0.0.30 sp2
System 3	10.0.0.32 sp3-firewall	10.0.0.33 sp3-ntc	10.0.0.40 sp3
System 4	10.0.0.42 sp4-firewall	10.0.0.43 sp4-ntc	10.0.0.50 sp4

TABLE 3-5 Naming Convention for Additional Systems (Continued)

	Firewall Address and Name	NTC Address and Name	Storage Service Processor Address and Name
System 5	10.0.0.52 sp5-firewall	10.0.0.53 sp5-ntc	10.0.0.60 sp5
System 6	10.0.0.62 sp6-firewall	10.0.0.63 sp6-ntc	10.0.0.70 sp6
System 7	10.0.0.72 sp7-firewall	10.0.0.73 sp7-ntc	10.0.0.80 sp7

About Network Configuration Options

The following table describes the choices to consider when deciding the level of security for the system.

TABLE 3-6 Network Configuration Considerations

Network Configuration	Description
Site local area network (LAN) connected to a system's firewall	Because the system is integrated into your network, you can log in to any host on the network and get access to the system to manage it. The system is visible to and can be discovered by third-party software. See "Managing From the Site LAN" on page 51 for definitions of the protocols and management options.
Site LAN with restricted services	Although the system is still visible on the network, access to it is restricted by the selective disabling of services through the system's firewall. For example, you can enable the management and diagnostic software but disable the ability of third-party software to discover the system. See "About Restricting Services" on page 52 for more information.
No site LAN; isolated system	To get access to the system, you must connect a computer to it. See "Administering a System Isolated From the LAN" on page 57 for more information.

Managing From the Site LAN

You can manage the storage in the system from any host that has a network connection to the system:

- You can administer the system using the system's user interface from any host in the network.
- You can manage the system and monitor, diagnose, and fix problems manually or by means of scripts using a remote scripting client. Remote scripting clients are available for the Solaris and Windows operating systems.
- You can monitor the system using third-party management tools such as Tivoli or BMC PATROL, which are supported by the Common Information Model (CIM), and SNMP interfaces to allow discovery of the system.

Setting Up Connections to the Site LAN

An Ethernet cable connects your site local area network (LAN) to the I/O section of the service panel. The cable runs from the panel's User LAN port either to a host or to the LAN's firewall or gateway. For more information on cable connections, see the *Sun StorEdge 6920 System Getting Started Guide*.

You have the following options for how you connect the system to the site LAN:

- Dynamic or static IP address. The system can get a new IP address from your site's Dynamic Host Configuration Protocol (DHCP) server each time it logs on to the LAN, or you can assign a specific IP address to it.
- The level of access that the system allows from the network.

About Data Host Connections

A data host is any host that uses the system for storage. The data host can be connected to the system in the following ways:

- To a host bus adapter (HBA), also called an initiator. The HBA is connected by a cable to a Fibre Channel (FC) port on the I/O section of the service panel.
- To an optional external switch, which allows fan-out from one of the FC ports to multiple hosts.

The number of hosts that can connect to the system without an external switch installed is limited by these factors:

- The number of storage resource card (SRC) sets installed in the base cabinet.
- The number of arrays installed.

If you connect an external switch to one or more of the system's FC ports, you increase the number of data hosts that can use the system.

About Restricting Services

If your site has a security requirement, you can block certain services on the system to control the degree of access from your site local area network (LAN). The following table shows which services can be blocked. The ports for Simple Mail Transfer Protocol (SMTP) (port 25), DNS (port 53), and Dynamic Host Configuration Protocol (DHCP) (port 67) are always open for outgoing messages.

TABLE 3-7 Restricting Services

Service	Allows Access To:	Port	Direction
Storage Management HTTP	Non-secure web or command-line interface (CLI)	9080	Incoming
Storage Management HTTPS	Secure web or CLI	6789	Incoming
Service HTTP	Non-secure web or CLI	7654	Incoming
Service HTTPS	Secure web or CLI	443	Incoming
Simple Network Management Protocol (SNMP)	Third-party tools	161	Incoming Outgoing
Common Information Model (CIM) HTTP	Third-party tools	5998	Incoming
SSH	Secure shell	22	Incoming
Revision Maintenance	Firmware or software updates	443	Outgoing

About Allocating Ports

The service panel includes two sections: the Service Processor panel on the top and the I/O panel on the bottom. The service panel provides the following connections:

- Modem connection
- Local area network (LAN) connections for management
- Serial ports
- Fibre Channel (FC) ports for data host, storage, external storage device, and Sun StorEdge 6920 system connections
- Gigabit Ethernet ports for connections to a remote Sun StorEdge 6920 system for replicating volume data between systems over a wide area network (WAN)
- Power connections for expansion cabinet management
- Ethernet and FC connections for expansion cabinet AUX ports to connect to the Data Services Platform (DSP) management interface card (MIC) card, which provides management interface failover.




FIGURE 3-1 Service Processor Panel

As a guideline, use the four leftmost FC ports, ports 1 to 4, to connect to data hosts, and use the four rightmost ports, ports 5 to 8, to connect to the arrays in the system. However, you can allocate the ports as needed, depending on the number of arrays.

For example, if your site has fewer than four arrays in a single-cabinet system connected to a storage resource card (SRC) set, the unused ports can be used for connecting more than four data hosts, subject to the following limitations:

- The total number of hosts that can be connected directly to each SRC set is limited by the number of initiators each host has and by the number of arrays. If you need to connect more than the maximum number of initiators that are supported by the number of SRC sets, you must add a front-end switch and connect the ports to the switches outside the system.
- The number and configuration of arrays that can be installed in expansion cabinets is limited by the number and configuration of arrays in the base cabinet and the number of FC ports dedicated to hosts.

About Storage Resource Card (SRC) Sets

Data hosts connect to storage resource card (SRC) sets through the Fibre Channel (FC) ports on the service panel. An SRC set consists of two cards: the SRC processor card and the storage I/O (SIO) Combo card. SRC sets are always configured as pairs, so each system has at least two SRC processor cards and two SIO Combo cards.

Each SRC set has the following characteristics:

- The cards' connectors fasten to the back of the service panel.
- The connection between the two cards is through the chassis midplane. The SRC processor card plugs into the front of the midplane, and the SIO Combo card plugs into the back of the midplane.
- The SIO Combo card provides the FC interface to the storage and the Gigabit Ethernet connection for data replication to a Sun StorEdge 6920 system at a remote site.
- You can swap either card in the set without affecting the other card.
- Each SRC set provides six FC ports and one Gigabit Ethernet port. The processors on the card recognize whether a host or storage device is attached to a port. Each SRC set usually has four FC ports allocated for storage and four FC ports allocated for data hosts.

Note: To make optimal use of the processors on the SRC sets, use the ports for storage connections in the following order: 8, 6, 7, 5, 4, 3, 2, 1.

The number of SRC sets determines the number of initiators that can be connected to the system. Each system can be configured with a minimum of two SRC sets for 16 FC ports, and a maximum of four SRC sets for 32 FC ports.

Guidelines for Fibre Channel Port Connections

When you add new hardware, such as a storage resource card (SRC) set, a high-availability (HA) storage tray configuration, a data host, or an external switch, verify the Fibre Channel (FC) port connections on the Service Processor panel and check how the ports are distributed between data hosts and storage. For more information, see “About Allocating Ports” on page 53.

In addition, consider the following:

- In general, use the four leftmost ports to connect to data hosts and the four rightmost ports to connect to the arrays in the storage system.
- In a direct-attached storage configuration, the total number of hosts that can be connected to each SRC set is limited by the number of initiators each host has and by the number of arrays. For example, if fewer than four arrays in a single-cabinet system are connected to an SRC, the unused storage ports can be used for connecting more than four initiators.
- No front-end switches are needed, unless you want to connect more than the maximum number of initiators that are supported by the number of SRCs installed. To connect more than the maximum supported initiators, connect the ports to FC switches outside the system for fan-out or fan-in.
- The number and configuration of arrays that can be installed in expansion cabinets is limited by the number of FC ports dedicated to hosts as well as by the number and configuration of arrays in the base cabinet.
- If you connect an external storage device to the system by means of a FC switch, you must apply zoning practices. For information about zoning practices, see “About Using External Storage” on page 191 and the FC switch vendor’s documentation.

About Host Bus Adapters

When a data host is connected to the Sun StorEdge 6920 system by a host bus adapter (HBA), the host bus adapter is also known as an initiator. The HBA is connected by a cable to a Fibre Channel (FC) port on the I/O section of the service panel.

After connecting a data host directly to the system with one or two HBAs, use the `luxadm(1M)` command to verify the firmware level of the HBAs. If the firmware is not at the correct revision, use the Sun StorEdge 6920 Host Installation Software CD to install the Sun StorEdge SAN Foundation software. You can then configure the initiator.

Connecting an External Switch

You have the option of adding an external switch to one or more of the data host ports to expand the number of data hosts that use the storage system and connect external storage devices. However, external switches are not managed by the system. They must be connected to an external local area network (LAN) and be managed from a management station attached to that LAN.

Before you add an external switch, do the following:

- Consider the performance implications of connecting additional data hosts to a single Fibre Channel (FC) port on the system. Performance is affected by the number of data hosts competing for access to storage through a single port.
- Review the list of supported switches.
- Download and install switch management software and configure the software and the switch.
- Install the Sun StorEdge SAN Foundation software and any needed patches related to the switch.

To install the external switch:

- a. Power off the system.
 - b. Connect an FC cable from an initiator port to the external switch port.
 - c. Connect an FC cable from the external switch port to the storage resource card (SRC) set on the I/O panel.
 - d. Use the switch's management software to place both ports on the external switch in the same zone.
 - e. Power on the system.
- Consider zoning practices, if you are using an external switch to connect external storage devices to the system. For more information about zoning practices, see "About Using External Storage" on page 191 and the FC switch vendor's documentation.

About Multipathing

With multipathing, also called multipath failover, a system or network can detect when an adapter has failed and automatically switch access to an alternate adapter. Multipathing enables high-availability configuration because it ensures that the data path remains active. Multipathing also helps increase performance to multi-controller disk arrays by spreading I/O between multiple paths into the array.

Within the system, storage pools use multipathing by default. To complete the data path, data hosts also need the ability to multipath. Therefore, all data hosts need one of the following software products:

- Sun StorEdge Traffic Manager, which is Sun's multipathing solution for Fibre Channel (FC)-connected storage devices. This software is provided as part of the Sun StorEdge SAN Foundation software. For more information about downloading, installing, and configuring this software, see the Sun StorEdge SAN Foundation software documentation.
- VERITAS Volume Manager with Dynamic Multipathing (DMP).

Administering a System Isolated From the LAN

Use the following procedure when your site's security requirements necessitate having the system isolated from any external local area network (LAN):

- Set up and configure the system using a management host that is connected by a serial cable to the system's service panel.
- When you have finished configuring the system, disconnect the management host.
- Whenever the system needs to be reconfigured, reconnect the management host.

Another way to secure the system is to install a firewall between the management host for the system and the external LAN. For more information, see the system configuration instructions in the *Sun StorEdge 6920 System Getting Started Guide*.

Installing the Remote Scripting CLI Client

If you need to configure the system from a management host, the Sun StorEdge 6920 Host Installation Software CD contains a remote scripting command-line interface (CLI) that enables you to do most of the tasks supported by the browser interface. The commands can be used in scripts you create, or they can be entered directly in a terminal window's command line. The client can run on the following types of hosts:

- Solaris
- IBM AIX
- Red Hat Linux
- HP-UX
- Windows 2000

To install the remote scripting CLI client on a Solaris management host:

1. Insert the Host Installation Software CD.
2. Run the `install.sh` script.
3. Select Remote Configuration CLI.

The command for management services is `sscs`, used with a subcommand to direct the operation. For a list of commands, see the `sscs(1M)` man page.

For more information, see the *Sun StorEdge 6920 System Getting Started Guide*.

About Lights-Out Management

Lights-out management, also called remote power management, enables you to use a remote console to power off and power on the system. The system has the lights-out management feature disabled by default so that the system's key switch controls power to the system components and the arrays. When you enable the lights-out management feature, the power to these components is controlled by the power relay's sequencers.

For instructions about enabling and disabling the lights-out management feature, follow these steps:

1. From the Java Web Console page, click Sun Storage Automated Diagnostic Environment.
2. Click Service > Service Advisor > X-Options.

Powering the System On and Off

By default, powering off the system results in a full shutdown. If the power relay is enabled, you have the option of performing a partial shutdown.

- In a partial shutdown, the system shuts down the storage devices and continues to run, enabling you to restart the system from a remote console. Because the Data Services Platform (DSP) is connected to switch outlets of the power sequencer, the system can restore power to it. Then the management interface card (MIC) powers on the other cards in the DSP.
- In a full shutdown, the system shuts down the storage devices, then gradually shuts down itself. You must restart the system manually.

You must be logged in as the `storage` user to power off the system. Before you power off the system, you must halt any I/O between the data hosts and the system.

Caution: Failure to stop I/O before powering off can cause data loss.

Performing a Partial Shutdown

Note: Before you power off the system, you must halt any I/O between host systems and the Sun StorEdge 6920 system.

When you perform a partial shutdown, the system is in the following state:

- The Storage Service Processor in the base cabinet remains powered on.
- The Storage Service Processor accessory tray remains powered on.
- All storage arrays are powered off.
- The Data Services Platform (DSP) is powered off.
- Only the Power Available light emitting diodes (LEDs) on the power sequencers are lit.

To perform a partial shutdown of the system:

1. Log in as a storage user.
2. Click Sun StorEdge 6920 Configuration Service > Administration > General Settings.

3. Click the System Partial Shutdown button.
4. Click OK to confirm you want to perform a partial shutdown.
The following message is displayed on the General Settings page:
Partial system shutdown completed successfully

Restoring the System After a Partial Shutdown

If the system has been powered off with the remote partial shutdown procedure, use the following procedure to restore power to the system:

1. Log in as a storage user.
2. Click Sun StorEdge 6920 Configuration Service > Administration > General Settings.
The General Settings page is displayed.
3. Click System Power Up.
This selection activates the power sequencers in the base cabinet and any expansion cabinet and causes the storage components to power on. The Switched Outlet light emitting diodes (LEDs) are illuminated.
The system is now ready for operation.

Performing a Full Shutdown

Note: Before you power off the system, you must halt any I/O between host systems and the Sun StorEdge 6920 system.

To turn the system's power off completely:

1. Log in as a storage user.
2. Click Sun StorEdge 6920 Configuration Service > Administration > General Settings.
The General Settings page is displayed.
3. Click the System Shutdown button.
4. Click OK to confirm you want to perform a complete shutdown.

The system is now in the following state:

- The Storage Service Processor in the base cabinet is powered off and under the control of lights-out management.

- The Storage Service Processor accessory tray remains powered on.
- The Data Services Platform (DSP) remains powered on.
- All storage trays remain powered on.
- The power and cooling unit (PCU) fans are on.
- The PCU Remove LEDs are lit.
- All Power LEDs on the power sequencers are lit.

You must perform the following manual procedure on the system to complete the shutdown:

1. Remove the front trim panel from the base cabinet and any expansion cabinets.
2. At the bottom front and bottom back of each cabinet, raise the AC power sequencer circuit breakers to Off.

The system is now in the following state:

- The Storage Service Processor accessory tray is powered off.
- The DSP is powered off.
- All storage trays are powered off.
- The PCU fans are off.

If you are servicing the power sequencers or moving the system, disconnect the power cables. Otherwise, leave the power cables connected to ensure a proper grounding path for electrostatic discharge.

Restoring the System After a Full Shutdown

If you want to restore the system after it has been powered off with the full shutdown procedure, you must go to the location of the system and perform the following procedure:

1. Open the front door and back door of the base cabinet and any expansion cabinets.
2. Remove the front trim panel from each cabinet.
3. Verify that the AC power cables are connected to the correct AC outlets.
4. At the bottom front and bottom back of each cabinet, lower the AC power sequencer circuit breakers to On.

The power status light emitting diodes (LEDs) on both the front and back panel illuminate in the following order, showing the status of the front power sequencer:

- The main AC power LED (this lights when the power is applied to the cabinet)
 - Sequencer Stage 1
 - Sequencer Stage 2
5. At the back of the system, locate the power switch for the Storage Service Processor and press the power switch on.
 6. Verify that all components have only green LEDs lit.
 7. Replace the front trim panels and close all doors.

The system is now operating and supports the remote power-on procedure.

Specifying General Administrative Settings

This section describes general administrative settings. It contains the following subsections:

- [“Adding a System Description” on page 63](#)
- [“Setting the System IP Address” on page 63](#)
- [“Enabling the NTP Server” on page 64](#)
- [“Setting the System Time” on page 65](#)
- [“Changing Passwords” on page 65](#)

Adding a System Description

Any user can view the information on this page; however, to add or modify the system description, you must be logged in as the admin or storage user.

To add a description for your system:

1. Click Sun StorEdge 6920 Configuration Service > Administration > General Setup.
2. Go to System Properties, and type up to 64 alphanumeric characters in the Description field.

The description can include underscores, dashes, periods, colons, commas, percent signs, parentheses, braces, brackets, ticks, tildes, bars, and spaces.

3. Click Save to apply your changes.

Setting the System IP Address

Any user can view the information on this page; however, to set or modify the Internet Protocol (IP) address, you must be logged in as the admin or storage user.

You must specify the method by which an IP address is supplied for the system:

- DHCP (Dynamic Host Configuration Protocol) – The network provides an IP address for the system each time the system is powered on and logs in to the network.

- Static – You must provide the IP address the system will use each time it is powered on and logs in to the network.

To set the system IP address:

1. Click Sun StorEdge 6920 Configuration Service > Administration > General Setup.
2. Go to Network Settings, and select either DHCP or Static in the IP assignment box.
3. If you selected DHCP, click Save.

If you selected Static, type the following additional information, and then click Save:

- IP address you are assigning to the system
- Gateway IP address you are assigning to the system
- Network Mask for the assigned gateway
- Name Server's IP address for the network's domain

Note: The web browser you use to manage your system relies on the IP address. If the IP address changes, your browser will lose its connection to the system. In this case, you must reconnect to the system to continue monitoring and managing the system.

Enabling the NTP Server

Any user can view the information on this page; however, to enable the Network Time Protocol (NTP) server, you must be logged in as the admin or storage user.

To configure your system's real-time clock to synchronize with your network's NTP server:

1. Click Sun StorEdge 6920 Configuration Service > Administration > General Setup.
2. Go to NTP, and select the check box to enable this option.
3. In the IP Address field, type the IP address for your network's NTP server.
4. Click Save to apply your changes.

Setting the System Time

Any user can view the information on this page; however, to set or modify the system time, you must be logged in as the admin or storage user.

If the system does not use your network's NTP server, you must manually set the system time.

To set the system time:

1. Click Sun StorEdge 6920 Configuration Service > Administration > General Setup.
2. Go to System Time, and make the appropriate selections for the following:
 - System Time (the current time)
 - Current Month
 - Current Day
 - Current Year
 - Time Zone
3. Click Save to apply your changes.

Changing Passwords

Any user can view the information on this page; however, to change the password for any account, you must be logged in as the admin or storage user.

Note: To change an admin or storage password that is lost or forgotten, you will need to contact Sun Services for assistance in restoring access to that user account, before it can be used.

To change a password:

1. Click Sun StorEdge 6920 Configuration Service > Administration > General Setup.
2. Go to Passwords.
3. From the User list, select the account for which you want to change the password.
4. In the New Password field, type the new password, consisting of up to eight alphanumeric characters.

Note: The new password cannot be identical to the user account name. For example, you cannot use “guest” as the password for the guest user account name.

5. Retype the new password in the Confirmation field.
6. Click Set Password to save your changes.

Performing Monitoring Administration

This section describes how to monitor and manage devices. It contains the following subsections:

- [“General Setup Tasks” on page 67](#)
- [“About Notification” on page 71](#)
- [“Email Configuration Tasks” on page 71](#)
- [“Using Email Filters” on page 74](#)
- [“Remote Notification Tasks” on page 76](#)
- [“Editing Notification Setup Parameters” on page 78](#)
- [“Agent Tasks” on page 78](#)
- [“Ejecting a CD-ROM” on page 80](#)

General Setup Tasks

Use general setup tasks to configure various system features.

This section describes general setup tasks. It contains the following subsections:

- [“Displaying and Editing Site Setup Information” on page 67](#)
- [“Displaying and Editing System Timeout Settings” on page 68](#)
- [“Displaying and Editing System Configuration Options” on page 68](#)
- [“Enabling/Disabling Fault Signature Analysis” on page 69](#)
- [“Configuring Parallel Monitoring Parameters” on page 69](#)
- [“Displaying and Editing MAC Addresses \(Ethers\)” on page 70](#)
- [“Exporting a Configuration” on page 70](#)

Displaying and Editing Site Setup Information

Site information parameters are required for service personnel to identify contact and installation information for the customer.

To display and edit site information:

1. Click Sun Storage Automated Diagnostic Environment > Administration.
The Site Setup page is displayed.
2. To edit any parameter:
 - a. Click in the value field for the parameter and enter the new value.
 - b. Click Save after completing all required changes.

Displaying and Editing System Timeout Settings

You can display and edit system timeout settings to match the specific requirements of your storage network. The default timeout settings values are appropriate for most storage area network (SAN) devices. However, network latencies, I/O loads, and other device and network characteristics may require that you customize these settings to meet your configuration requirements.

To display and edit system timeout settings:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Timeout Settings.
The Timeout Settings page is displayed.
2. To edit any parameter:
 - a. Click in the value field for the parameter and enter the new value.
 - b. When all required changes are complete, click Save.

Displaying and Editing System Configuration Options

To display and edit general setup parameters:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Options.
The Options page is displayed.
2. To edit any parameter in the System Configuration Options section:
 - a. Click in the value field for the parameter and enter the new value.
 - b. Click Save after completing all required changes.

Enabling/Disabling Fault Signature Analysis

When you enable fault signature analysis, the system groups together events that share common field-replaceable units (FRUs). This reduces the number of alarms associated with an FRU problem.

To enable fault signature analysis:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Options.

The Options page is displayed.

2. In the Fault Signature Analysis section, click Yes.
3. Click Save.

To disable fault signature analysis:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Options.

The Options page is displayed.

2. In the Fault Signature Analysis section, click No.
3. Click Save.

Configuring Parallel Monitoring Parameters

You can specify the number of switches and Sun StorEdge arrays to be probed simultaneously. Probing more than one device at a time accelerates the monitoring frequency of the device type and improves the timeliness of event reporting.

Note: Increasing the number of devices that are monitored in parallel increases the amount of system resources that are used. It is recommended that no more than 5-10 devices be monitored in parallel.

To configure parallel monitoring parameters:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Options.

The Options page is displayed.

2. From the Number of switches to probe at the same time list, select the maximum number of switches to be probed simultaneously

3. From the Number of Sun StorEdge arrays to probe at the same time list, select the maximum number of arrays to be probed simultaneously
4. Click Save.

Displaying and Editing MAC Addresses (Ethers)

Ethers enables you to associate a Media Access Control (MAC) address with each device Internet Protocol (IP) name or IP address. Each MAC address identifies a storage location or device for each IP name or IP address. When you enter or update a MAC address for an array, the array can be automatically assigned an IP address at boot time.

Note: An Ethernet address is assigned at the factory to the Ethernet controller board. Network communication requires that the fixed Ethernet address be mapped to the IP address of the device.

To associate a Media Access Control (MAC) address with a device:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Ethers.
The Ethers page is displayed.
2. Add, change, or delete MAC addresses for listed devices.
Note: MAC addresses for each device are found on a sticker attached to the back panel of the device.
3. Click Save to update the file with the specified changes.
4. Turn power off and on for each device for which you have newly assigned a MAC address.

Exporting a Configuration

You can export an HTML-formatted report that details the system environment and configuration and email it to a specified recipient.

To export configuration information:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Utilities > Export Configuration.
The Export Configuration page is displayed.
2. In the Enter Target Email field, specify the email address to which you want to export the configuration information.

3. (Optional) Specify a text message in the Message field.
4. Click Submit.

Detailed information is sent to the specified email address.

About Notification

The monitoring and diagnostic software provides mechanisms to provide alarm notification to local and remote recipients.

- **Email notification** – Email notification options include the ability to send alarm notification to one or more email and pager addresses. For each recipient, you can define the level of alarms that will trigger notification.
- **Remote notification** – Remote notification options include:
 - **Simple Network Management Protocol (SNMP) traps** – You can configure remote notification to send SNMP traps upon receipt of alarms. You can also configure the minimum alarm level at which SNMP traps are sent.
 - **Sun StorEdge Remote Response service** – You can configure remote notification to send alarm notification to the Sun StorEdge Remote Response service, a Sun Solution service that provides dial-up connectivity from the Sun StorEdge 6920 array. The service, labeled SSRR in the browser interface, provides remote response to alarms and updates system firmware.
 - **Network Storage Command Center (NSCC)** – You can configure remote notification to send alarm information to Sun’s NSCC, a database used to collect data on Sun storage device health and performance. The data is used by Sun to improve system performance. You can configure NSCC notification to encrypt all transmitted data.

Email Configuration Tasks

This section describes email configuration tasks. It contains the following subsections:

- [“Configuring Email Notification” on page 72](#)
- [“Testing a Notification Email Address” on page 72](#)
- [“Adding an Email Notification Recipient” on page 73](#)
- [“Editing an Email Notification Recipient” on page 73](#)
- [“Deleting an Email Notification Recipient” on page 74](#)

Configuring Email Notification

To configure email notification:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification.

The Notification Setup page displays notification setup parameters.

2. If the Sun StorEdge 6920 system is installed behind a firewall, specify an Simple Mail Transfer Protocol (SMTP) server to use for email instead of the host server.
3. Specify a server path to email program if an SMTP server is unavailable.
4. (Optional) Specify an email address to be used as the sender's address for all sent email.
5. Specify the maximum size allowed for email messages to be sent.
6. Click Save.

The configured email notification setup parameters are saved.

Testing a Notification Email Address

To test an email notification address:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification.

The Notification Setup page displays notification setup parameters.

2. Specify the Use this Simple Mail Transfer Protocol (SMTP) server for Email parameter.
3. Click Test Email.

The Send Test Email window is displayed.

4. In the To field, specify the email address that you want to test.
5. Optionally, specify a text message in the Message field.
6. Click Test Email.

A confirmation message appears at the top of the Send Test Email window if the test email is successfully sent.

Adding an Email Notification Recipient

To add an email notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email.

The Email Notification page is displayed.

2. Click New.

The Add Email Notification page is displayed.

3. Define the email address and, optionally, other email properties for the notification recipient.

4. Click Save.

The Email Notification page is redisplayed, and the new email notification recipient is listed.

Editing an Email Notification Recipient

To edit email properties for an email notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email.

The Email Notification page is displayed.

2. Select the check box of the email notification recipient that you want to edit.

3. Click Edit.

The Edit Email Notification page is displayed.

4. Edit the email properties for the notification recipient as needed.

5. Click Save.

The Email Notification page is redisplayed, and the edited email notification recipient is listed. The following confirmation message is displayed at the top of the page: `This notification successfully updated.`

Deleting an Email Notification Recipient

To delete an email notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email.

The Email Notification page is displayed.

2. Select the check box to the left of each email notification recipient that you want to delete.
3. Click Delete.

The Email Notification page is redisplayed, and the selected email notification recipients are no longer listed.

Using Email Filters

This section describes email filters. It contains the following subsections:

- [“About Email Filters” on page 74](#)
- [“Adding an Email Filter” on page 75](#)
- [“Editing an Email Filter” on page 75](#)
- [“Deleting an Email Filter” on page 75](#)

About Email Filters

You can use email filters to prevent email notification about specific events.

To prevent email notification about specific events that occur frequently, obtain the event code for the event from the Event Details page of the event and add an email filter to prevent email notification for events with that event code.

You can still view filtered events in the event log.

Adding an Email Filter

To add an email filter:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email Filters.
The Email Filters page is displayed.
2. Click Add New Filter.
The Add New Filter page is displayed.
3. Define the Event Code and Decreased Severity for the filter, and click Save.
The Email Filters page is redisplayed with the new email filter.

Editing an Email Filter

To edit an email filter:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email Filters.
The Email Filters page is displayed.
2. Select the check box of the filter you want to edit and click Edit.
The Edit Email Filter page is displayed.
3. Edit the Event Code and Decreased Severity parameters as needed, and click Save.
The Email Filters page is redisplayed with the edited email filter.

Deleting an Email Filter

To delete an email filter:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > Email Filters.
The Email Filters page is displayed.
2. Select the check boxes of the filters you want to delete, and click Delete.
The Email Filters page is redisplayed without the deleted email filters.

Remote Notification Tasks

This section describes remote notification tasks. It contains the following subsections:

- [“Configuring Remote Notification” on page 76](#)
- [“Adding an SNMP Notification Recipient” on page 76](#)
- [“Editing an SNMP Notification Recipient” on page 77](#)
- [“Deleting an SNMP Notification Recipient” on page 77](#)

Configuring Remote Notification

To configure notification setup parameters:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification.

The Notification Setup page displays notification setup parameters.

2. Select the check box of each remote notification provider you want to enable.
3. If Network Storage Command Center (NSCC) notification is enabled, specify whether notification email to the NSCC should be encrypted.
4. If Sun StorEdge Remote Response (SSRR) notification is enabled, specify a Heartbeat Frequency for the service.
5. If Simple Network Management Protocol (SNMP) Traps is enabled, specify one or more SNMP trap recipients.
6. Click Save.

The configured remote notification setup parameters are saved.

Adding an SNMP Notification Recipient

To add a Simple Network Management Protocol (SNMP) notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > SNMP.

The SNMP Notification page is displayed.

2. Click New.

The Add SNMP Notification page is displayed.

3. Define the SNMP Properties for the new SNMP notification recipient.

4. Click OK.

The SNMP Notification page is redisplayed, and the new SNMP notification recipient is listed.

Editing an SNMP Notification Recipient

To edit Simple Network Management Protocol (SNMP) properties for an SNMP notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > SNMP.

The SNMP Notification page is displayed.

2. Select the check box of the SNMP notification recipient that you want to edit.

3. Click Edit.

The Edit SNMP Notification page is displayed.

4. Edit the SNMP Properties for the notification recipient as needed.

5. Click OK.

The SNMP Notification page is redisplayed and the edited SNMP notification recipient is listed. The confirmation message, "This notification successfully updated," is displayed at the top of the page.

Deleting an SNMP Notification Recipient

To delete an Simple Network Management Protocol (SNMP) notification recipient:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification > SNMP.

The SNMP Notification page is displayed.

2. Select the check box to the left of each SNMP notification recipient that you want to delete.

3. Click Delete.

The SNMP Notification page is redisplayed and the selected SNMP notification recipients are no longer listed.

Editing Notification Setup Parameters

To reset notification setup parameters to their last saved values:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification.
The Notification Setup page displays notification setup parameters.
 2. Configure Email Notification and Remote Notification setup parameters as needed.
 3. (Optional) Click Reset to return all notification setup parameters to their last saved values.
-

Agent Tasks

This section describes agent tasks. It contains the following subsections:

- [“Activating/Deactivating the Agent” on page 78](#)
- [“Running the Agent” on page 79](#)
- [“Displaying Agent Statistics” on page 79](#)
- [“Clearing Agent Statistics” on page 79](#)

Activating/Deactivating the Agent

To activate or deactivate the agent:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Agent.
The Agent Summary page is displayed.
2. Click Deactivate to deactivate the agent or Activate to activate the agent.
A confirmation message is displayed.

Running the Agent

To run the agent manually:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Agent.
The Agent Summary page is displayed.
2. (Optional) To audit all devices, select Audit All Devices.
3. (Optional) To ignore email maximums when running the agent, select Ignore Email Maximums.
4. Click Run.
The Run Agent Results page is displayed.

Displaying Agent Statistics

To display agent statistics, click File System Manager > Administration > Agent > Statistics.

The Agent Statistics page is displayed.

Clearing Agent Statistics

To clear agent statistics:

1. Click File System Manager > Administration > Agents > Statistics.
The Agent Statistics page is displayed.
2. Click Clear Statistics.
The Agent Statistics page is redisplayed with all statistics cleared.

Ejecting a CD-ROM

The Eject CD utility ejects a compact disc read-only memory (CD-ROM) disc from the Storage Service Processor.

To eject the CD-ROM:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Utilities > Eject CD.

The Eject CD page is displayed.

2. Click Eject.

The Storage Service Processor CD is ejected.

Performing Configuration and Management Tasks

This chapter describes configuration and management tasks. It contains the following sections:

- “Configuring Feature Licensing” on page 82
- “Configuring Storage Volumes” on page 87
- “Configuring Mirrored Volumes” on page 93
- “Configuring Snapshots” on page 103
- “Using Data Replication” on page 116
- “Configuring Virtual Disks” on page 150
- “Configuring Storage Pools” on page 153
- “Configuring Storage Profiles” on page 158
- “Configuring Storage Domains” on page 167
- “Configuring Initiators” on page 171
- “Configuring System Ports” on page 175
- “Configuring Storage Arrays” on page 180
- “Using Storage Trays” on page 184
- “Managing Disk Drives” on page 189
- “Managing External Storage” on page 191

Configuring Feature Licensing

This section describes licensing features that are available on your system. It contains the following subsections:

- [“About Feature Licenses” on page 82](#)
- [“Displaying Feature License Compliance” on page 84](#)
- [“Adding a Feature License Key” on page 84](#)
- [“Displaying a Feature License Key” on page 85](#)
- [“Deleting a Feature License Key” on page 85](#)

About Feature Licenses

Before you can use premium data services features on your system, you must obtain a feature license key and register feature licenses for each feature that you plan to use.

Note: Though you can exceed the licensed capacity of a data service, the non-observance of license requirements is visible to Sun personnel who monitor the Network Storage Command Center (NSCC).

When you purchase a feature license, you receive a feature license key that entitles you to use that feature. These feature licenses are capacity-based licenses. If necessary, you can increase the capacity of the feature license at a later time. Contact your Sun sales representative for more information about purchasing feature licenses.

Note: Your system might include factory-installed feature licenses.

The features that you must license include the following:

- Sun StorEdge Data Mirror software

The feature license grants you the right to use a specific amount of storage on a volume, and use that volume to create a mirror.

If you create a mirror from an existing volume, the feature license pertains only to the original volume’s storage amount and not to the number of components that make up the mirror. For example, an original volume that is licensed for 1 Gbyte of storage can be part of a two-component mirror or a four-component mirror, even though the latter requires more consumed storage to create the mirror.

If you create a new mirrored volume, you must specify the storage amount of the mirror. The feature license pertains to the storage amount you specify.

Note: If you increase the size of an existing mirror, you are required to increase the feature license amount.

For more information about the data mirroring feature, see “Configuring Mirrored Volumes” on page 93.

- Sun StorEdge Data Replicator software for asynchronous replication

The feature license grants you the right to use a specific amount of storage on a volume, and use that volume to replicate data to a remote site using asynchronous communication. You must register a feature license on the local system and a separate feature license on the system at the remote site to which you plan to replicate volumes.

For more information about the data replication feature, see “Managing Data Replication” on page 139.

- Sun StorEdge Data Replicator software for synchronous replication

The feature license grants you the right to use a specific amount of storage on a volume, and use that volume to replicate data to a remote site using synchronous communication. You must register a feature license on the local system and a separate feature license on the system at the remote site to which you plan to replicate volumes.

For more information about the data replication feature, see “Managing Data Replication” on page 139.

- Sun StorEdge Data Snapshot software

The feature license grants you the right to use a specific amount of storage on a volume, and use that volume to store point-in-time copies, or snapshots, of volumes of storage. The feature license pertains only to the volume’s storage amount, and not to the number of snapshots, the size of the snapshots, and the amount of storage consumed by the snapshot reserve space. For example, a 10-Gbyte volume requires a 10-Gbyte feature license, whether you create no snapshots or five snapshots. That same volume requires a 10-Gbyte feature license whether you configure a small or large snapshot reserve space.

For more information about the data snapshot feature, see “Configuring Snapshots” on page 103.

- Sun StorEdge Storage Pool Manager software for Sun storage

The feature license grants you the right to use a specific amount of internal or external Sun storage on the system. If you designate storage on an external Sun storage device, either as raw storage that you place into storage pools on the system or as legacy volumes that are available for use by the system, this external Sun storage is also part of the licensed amount of storage.

For more information about storage pool management, see “Configuring Storage Pools” on page 153.

- Sun StorEdge Storage Pool Manager software for non-Sun storage

The feature license grants you the right to use a specific amount of storage from an external non-Sun storage device that is connected to the system. If you designate storage on an external non-Sun storage device, either as raw storage that you place into storage pools on the system or as legacy volumes that are available for use by the system, this external non-Sun storage is part of the licensed amount of storage.

For more information about the external storage feature, see “Managing External Storage” on page 191.

If you use a feature but do not register any feature licenses, or if you exceed the capacity of the sum of the feature license keys for a feature, the Feature License Summary page indicates that your system is not in compliance with licensing requirements. You can determine the features for which you must contact your Sun sales representative to purchase additional feature licenses, and enter enough feature license keys to ensure that your system is in compliance with licensing requirements.

Note: The system saves feature license keys in the USB flash disk. This ensures that feature license keys remain valid in the event of a change to the system, such as installing a new Storage Service Processor.

Displaying Feature License Compliance

To display summary and detailed information about features that require feature licenses, and to audit feature license compliance of your system:

1. Click Sun StorEdge 6920 Configuration Service > Administration > Licensing.

The Feature License Summary page is displayed.

2. Click the name of the features for which you want to see feature license information.

The License Details page for the selected feature is displayed.

Adding a Feature License Key

To add feature license keys, you must be logged in to the system using the storage or administrator user account. Refer to “About User Roles” on page 18 for more information about user accounts.

To add a feature license key and register a feature for use:

1. Click Sun StorEdge 6920 Configuration Service > Administration > Licensing.
The Feature License Summary page is displayed.
2. Click Add.
The Add Feature License page is displayed.
3. In the Key field, enter the feature license key for the feature that you want to register.
4. Click OK.
The feature is registered for use.
5. Refer to the Feature License Summary page to see the total licensed storage amount and the number of licenses for the feature.

Note: If the system rejects a feature license key as invalid, make sure that you did not enter the feature license key incorrectly or did not enter a feature license key that is already in use.

Displaying a Feature License Key

To display a feature license key for a registered feature:

1. Click Sun StorEdge 6920 Configuration Service > Administration > Licensing.
The Feature License Summary page is displayed.
2. Click the name of the feature whose feature license key you want to see.
The License Details page for the selected feature is displayed.
3. See the License Key field for the feature license key.

Deleting a Feature License Key

You can transfer a feature license to another Sun StorEdge 6920 system. You must first delete the feature license on your system before you can use it on another system. To delete feature license keys, you must be logged in to the system using the `storage` or `administrator` user account. Refer to “About User Roles” on page 18 for more information about user accounts.

To delete a feature license key so that a feature is no longer registered for use:

1. Click Sun StorEdge 6920 Configuration Service > Administration > Licensing.
The Feature License Summary page is displayed.
2. Click the name of the feature for which you want to remove the feature license key.
The License Details page for the selected feature is displayed.
3. In the Actions column, click Delete.
The system prompts you to confirm that you want to remove the feature license key.
4. Click OK.
The feature is no longer registered for use.
5. Delete any volumes that used the feature to remain in compliance with licensing requirements.

Configuring Storage Volumes

This section describes information about storage volumes. It contains the following subsections:

- [“About Volumes” on page 87](#)
- [“Planning Volumes” on page 88](#)
- [“Managing Volumes” on page 88](#)

About Volumes

A volume is a logically contiguous range of storage blocks allocated from a single pool and presented by a disk array as a logical unit number (LUN). A volume can span the physical devices composing the array or be wholly contained within a single physical disk, depending on its virtualization strategy, size, and the array internal configuration. The array controller makes these details transparent to applications running on the attached server system. The Sun StorEdge 6920 system presents a volume as a logical entity that can be mapped to a host and allocated from a storage pool.

You can create volumes to efficiently manage data for your storage system. Volumes act as containers into which applications, databases, and file systems can store data.

You manage the system’s physical storage by placing virtual disks in a pool of storage space for creating volumes. Volumes are created from the virtual disks in storage pools. Based on your specifications, the system automatically allocates the storage from the appropriate virtual disks to satisfy your volume configuration requirements.

After you create volumes, assign each initiator in your environment to one of the volumes. The system provides wizards to assist you in creating volumes, creating storage pools, and assigning initiators.

While you should make every attempt to appropriately configure the capacity for volumes you create on your system, you can extend volumes when you find their configured capacity is not sufficient. If you choose to extend a volume, be sure to back it up first in order to preserve your data.

To see the current volumes, go to the Volume Summary page, as described in “Displaying Volume Information” on page 89.

Planning Volumes

You manage the system's physical disks as pools of storage space for creating volumes. Volumes are created from the virtual disks in storage pools. Based on your specifications, the system automatically allocates the storage from different disks to satisfy your volume configuration requirements.

When you create a volume, the system allocates capacity from a storage pool. You will most likely also map the volume to initiators. You can create volume snapshots, and you can extend the volume capacity if necessary.

To create a new volume, you need to know the following:

- The storage domain that will be associated with the volume
- The volume type – single or mirrored
- The storage pool in which the volume will be created
- The capacity for the volume – minimum 16 Mbytes; maximum 2 Tbytes
- The snapshot reserve space needed if snapshots will be created for the volume
- The initiators you want to map to the volume

Managing Volumes

This section describes the ways in which you can manage volumes. It contains the following subsections:

- [“Displaying Volume Information” on page 89](#)
- [“Creating a Volume” on page 89](#)
- [“Modifying a Volume” on page 90](#)
- [“Deleting a Volume” on page 90](#)
- [“Mapping a Volume to Initiators” on page 90](#)
- [“Unmapping a Volume From Initiators” on page 91](#)
- [“Extending a Volume” on page 91](#)
- [“Extending Replicated Volumes” on page 92](#)

Displaying Volume Information

You can display summary information and details about existing storage volumes. You can also view additional information about initiators, snapshots, and virtual disks associated with each volume.

To display information about volumes:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select a volume to see its detailed information.
The Volume Details page is displayed.
3. Go to Additional Information to see the number of mapped initiators, replication sets, snapshots, and virtual disks associated with the volume.
Click an item to see its Summary page.

Creating a Volume

You must consider a number of factors and make a number of decisions before creating a volume. For information on planning a volume, see “Planning Volumes” on page 88.

Note: Each volume must have a unique name, regardless of the domain in which it resides.

To create a volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Click New.
The New Volume wizard is displayed.
3. Follow the steps in the wizard.
Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Modifying a Volume

Note: Each volume must have a unique name, regardless of the domain in which it resides.

To modify a volume's name or description:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select the volume that you want to modify.
The Volume Details page is displayed.
3. Make the appropriate modifications, and click Save.

Deleting a Volume

You cannot delete a volume that is mapped to an initiator or that has snapshots associated with it. If the volume you want to delete is mapped to initiators, see "Unmapping a Volume From Initiators" on page 91 before following this procedure.

To delete a volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select the check box for the volume that you want to delete.
This enables the Delete button.
3. Click Delete.
The volume is removed from the Volumes Summary page.

Mapping a Volume to Initiators

To map a volume to one or more initiators:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select the check box for the volume to which you want to map one or more initiators.
This enables the Map button.

3. Click Map.

The Map Volumes wizard is displayed.

4. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Unmapping a Volume From Initiators

Before you can delete a volume that is mapped to an initiator, you must unmap it.

To unmap a volume from one or more initiators:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Click Mapped in the State column for the volume you want to unmap.

The Initiator Mapping Summary page appears.

3. Select the check box for the initiators from which you want to unmap the volume.

This enables the Unmap button.

4. Click Unmap.

5. Click OK to confirm the unmapping of the selected initiators.

Upon completion of the unmapping process, a confirmation message is displayed.

Extending a Volume

Before you extend a volume, be aware of the following:

- You cannot extend the capacity of a volume that has snapshots associated with it. You must first delete the snapshots.
- When you extend one mirrored volume, the capacities of all other components in that mirror are equally increased, regardless of whether they are on different arrays.

Warning: Be sure to fully back up a volume before you follow this procedure, to ensure your data is preserved if issues occur when you attempt to extend it.

To extend the capacity of a volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select the volume whose capacity you want to extend.
The Volume Details page is displayed.
3. Click Extend.
The Extend Volume page is displayed.
4. Select Maximum Capacity or specify the new capacity you want.
Note: The maximum capacity for a volume is 2 Tbytes.
5. Click OK.

Extending Replicated Volumes

You must consider a number of factors and make a number of decisions before extending replicated volumes. For information on planning replicated volumes, see “Planning for Data Replication” on page 134.

After you have created a replication set, you can perform typical volume operations.

To extend replicated volumes:

1. Place the replication sets on the primary and secondary peers in suspended mode.
2. Use the extend volume operation to extend the storage of the secondary volume.
3. Use the extend volume operation to extend the storage of the primary volume.
4. Initiate the resume operation for the replication set to synchronize both peers and resume replication.

The system displays an error message if the secondary volume is smaller than the primary volume.

Note: Do not extend a legacy volume if it is the primary or secondary volume in a replication set. If a legacy volume is a secondary volume in a replication set, do not extend the primary volume beyond the size of the legacy volume.

Configuring Mirrored Volumes

This section describes mirrored volumes. It contains the following subsections:

- [“About Mirroring” on page 93](#)
- [“Planning Mirror Components” on page 95](#)
- [“Managing Mirror Components” on page 97](#)

About Mirroring

Mirroring is a form of storage whereby two or more identical copies of data are maintained on separate media. Also called RAID Level 1 and real-time copy, typical mirroring technologies enable the cloning of data sets to provide redundancy for your system.

Surpassing industry standards, the Sun StorEdge 6920 system empowers you to create and manipulate a mirror with up to four equal mirror components that you can access independently. You can also update each independent copy of the mirror’s data set by performing a variety of operations on any mirror component within the mirror.

The Sun StorEdge 6920 system treats a mirror as a special type of volume made up of mirror components that you can use to track changes to your data set, update your data set, and manage your data migration strategy.

A mirror is made up of equal, independent mirror components periodically synchronized (through the resilvering process) to reflect the same data set. Although the first mirrored volume (also called the original volume and the primary component) you create establishes the initial data set for the mirror that will be copied to mirror components you add, the various operations that you perform on the mirror components affect the entire mirror. You can use any mirror component for any of the following: as a picture of your data at a specific time, to update the data set for the mirror, or for data-redundancy protection.

Note: The mirror assumes the World Wide Name (WWN) of the first mirrored volume you create to establish the mirror. Any existing volumes you add to the mirror preserve the WWN of the first mirrored volume you create.

Mirroring can be an important part of your data migration strategy. Suppose you want to migrate data from legacy to regular storage. You can mirror the data on a legacy volume to a normal volume, during which time the system can continue to access the legacy volume. When mirroring is complete, the mirror components have

been resilvered so that the legacy data also exists as an independent copy on the new mirrored volume. Since you can access your data on the new mirrored volume, you may choose to delete the legacy volume.

You can perform the following actions on mirror components within a mirror:

- **Split** – You can separate a mirror component from the mirror, with the intent to rejoin it later. Because a split mirror component continues to be tracked and maintains an association with other components within the mirror, it counts toward the limit of four mirror components per mirror. You can use a split mirror component to create a picture of your data at the time of the split. Because mirror components are independent copies of the same data set, a split mirror component offers a picture of data that, unlike a snapshot, is not dependent upon snapshot reserve space.
- **Rejoin** – You can join a split mirror component back into the mirror. Resilvering immediately begins, and the previously split mirror component becomes identical to (is synchronized to match) its counterparts in the mirror.
- **Reverse rejoin** – You can perform a reverse rejoin action on a split mirror component when you want to use its data as the basis for the mirror as a whole. The result is that every mirror component in the mirror becomes identical to the previously split mirror component, and the data set represented in the mirror matches that of the previously split mirror component.
- **Break** – You can remove a mirror component from the mirror, voiding its association with the other mirror components. The broken mirror component becomes a standalone volume in the system, and synchronization with the other mirror components is not maintained.

You might break a mirror component in order to create an independent copy of an existing volume while the mirror is active and save it as a picture of your data at a specific time.

Note: If you choose to add a previously broken mirror component to the mirror, you will add it as a normal volume. This is different from the split operation, because any changes in data for a split mirror component continue to be tracked, and a relationship between the split component and other mirror components is maintained. Typically, you will break a mirror component when your intent is to remove it from the mirror, and you will split a mirror component when your intent is to separate it and rejoin it later.

Note: If you break the last remaining component of a mirror, the result will be a non-mirrored volume with the same name as the mirror. The name of the previous mirror component is no longer used.

- **Force break** – You can execute a force break operation to remove a mirror component that has not been 100% resilvered and is therefore not synchronized with the other components in the mirror.

If you know that a mirror component's data has been compromised (during a fire or other catastrophe, for instance) and data integrity is lost, you can force the break of a mirror component, regardless of the percentage of resilvering that has occurred, to protect the integrity of the mirror's data. Just as with a break operation, the mirror component for which you implement a force break becomes a standalone volume in the system; however, the condition of its data is degraded, and it is inaccessible after being removed from the mirror during the resilvering process.

Planning Mirror Components

A mirror contains up to four independent, identical copies of your data in separate yet equal mirror components you access to manage your mirror as a whole. Mirror components are individual parts of the mirror, identical in both data and capacity. An action you perform on one mirror component can cause all other mirror components to synchronize to it, and vice versa.

For example, when you rejoin a split mirror component back into the mirror, it is synchronized (resilvered) to reflect the other mirror components. When the resilvering process completes, the mirror is once again an entity in which each component is separate subentity with identical data.

To create a mirror, you can create a new mirrored volume or mirror an existing volume and then add other volumes as equal mirror components.

Before starting the process, you must determine the following:

- The mirror components that will make up the mirror. Keep in mind the following restrictions:
 - Because mirror components must be identical in size, you must ensure each mirror component you add is the same size as the other components in the mirror.
 - You cannot mirror across storage domains.
- The capacity for the original mirrored volume (primary component). This capacity will be the same for new mirror components that are added to the mirror.
- The isolation policy you want for mirror components:
 - Required – Mirror components must reside on different arrays.
When mirror components exist on different arrays, redundancy protection is at its highest for data preservation.
 - Optional – Mirror components can reside on the same array.

If you know that a mirror component already shares an array with parts of your mirror, you must set the isolation policy to Optional.

When mirror components share an array, the redundancy factor is reduced for the data set associated with each mirror component, should the array or access to the array fail. If this possibility poses an issue, it might be safer to select the Required option.

- The resilver priority setting you want. Resilvering synchronizes each component to be identical with the others. It occurs each time you perform an operation on the mirror. This setting determines whether I/O activity or resilvering is a higher priority for the system:
 - High – The system attempts to resilver the mirror quickly, but this may noticeably slow down host I/O activities.
 - Medium – The system balances I/O activities with the resilvering process, yielding an incremental, steady pace for all jobs being executed.
 - Low – The system takes longer to resilver the mirror, but it processes host I/O activities more quickly.
- Whether you want to use split mirror components or volume snapshots for pictures of your data at a specific time:
 - Split mirror components – When you want to take a picture of the data set represented in a mirror, you can split a mirror component from the mirror and save it as an independent copy or reference point of your data at that time. Although you cannot roll back mirrored volume snapshots (because their data is unchangeable), you can split a mirror component and later rejoin it to the mirror, and you can even perform a reverse rejoin action on the previously split component if you want the other mirror components to synchronize to its data. The split mirror component has no snapshot reserve space limitations.
 - Volume snapshots – The volume snapshot feature is most useful as a method of tracking data through precise, point-in-time pictures that reside in snapshot reserve space. You can modify the snapshot reserve space of a mirrored volume as you would any other volume. The older a volume snapshot, the more the data it represents is likely to have changed. If the volume is expected to change substantially between snapshots, using split mirror components may be more appropriate. For more information, see “About Snapshots” on page 103 and “Managing Snapshots” on page 111.
- If the selected mirror components do not have snapshot reserve space and you want to take snapshots, how to allocate snapshot reserve space. For more information, see “Planning Snapshots and Snapshot Reserve Space” on page 106.

Note: The mirror assumes the World Wide Name (WWN) of the first mirrored volume you create to establish the mirror. This unique identifier is a 64-bit number that identifies the naming authority, manufacturer, and specific connection. Any existing volumes you add to the mirror preserve the WWN of the first mirrored volume you create.

Managing Mirror Components

This section describes ways in which you can manage mirror components. It contains the following subsections:

- [“Displaying Mirror Component Information” on page 97](#)
- [“Creating a Mirrored Volume” on page 97](#)
- [“Mirroring an Existing Volume” on page 98](#)
- [“Adding a Mirror Component” on page 99](#)
- [“Splitting a Mirror Component” on page 99](#)
- [“Rejoining a Split Mirror Component” on page 100](#)
- [“Reverse Rejoining a Split Mirror Component” on page 101](#)
- [“Breaking a Mirror Component” on page 101](#)
- [“Forcing the Break of a Mirror Component” on page 102](#)

Displaying Mirror Component Information

To display information about existing mirror components:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select a mirrored volume.
The Mirrored Volume Details page for the selected mirrored volume is displayed.
3. In the Mirror section of the page, select one of the components to see its details.
The Mirror Component Details page is displayed.

Creating a Mirrored Volume

You must consider a number of factors and make a number of decisions before creating a mirrored volume. For information on planning a mirrored volume, see [“Planning Mirror Components” on page 95](#).

To create a new mirrored volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.

2. Click New.

The New Mirrored Volume wizard is displayed.

3. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: The mirror assumes the World Wide Name (WWN) of the first mirrored volume you create to establish the mirror. The WWN is a 64-bit number that provides a unique identifier for the naming authority, manufacturer, and specific connection. Any existing volumes you add to the mirror preserve the WWN of the first mirrored volume you create.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Mirroring an Existing Volume

You must consider a number of factors and make a number of decisions before creating a mirrored volume from a volume that currently exists in the system. For information on planning a mirrored volume, see “Planning Mirror Components” on page 95.

To mirror an existing volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Select the volume you want to mirror. Its type must be Single.

The Volume Details page is displayed.

3. Click Mirror.

The Mirrored Volume wizard is displayed.

4. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Adding a Mirror Component

You must consider a number of factors and make a number of decisions before adding a mirror component to a mirror. For information on planning mirror components, see “Planning Mirror Components” on page 95.

To add a mirror component to a mirror:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select a mirrored volume that is part of the mirror to which you want to add a component.
The Mirrored Volume Details page is displayed.
3. In the Mirror section of the page, click Add.
The Add To Mirrored Volume wizard is displayed.
4. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

The Mirrored Volume Details page is displayed, with the Mirror section updated to show that the new mirror component has been added to the mirror. The component’s condition is listed as OK, Resilvering In Progress until the data set of the selected component is 100% resilvered and identical to the data sets of the other components in the mirror.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Splitting a Mirror Component

You can perform a split operation to temporarily separate a mirror component from the mirror. The system will persistently record changes you make to the split component, as well as changes made to the mirror; thus, it counts toward the four components allowed in a mirror. You can view the split mirror component as a picture of the data for the mirror at the time of the split. In this way, a split mirror component is similar to a volume snapshot; however, a split mirror component acts as an independent copy of the data set represented in the mirror, whereas a snapshot is a copy that is dependent upon snapshot reserve space.

Note: You cannot split a mirror component that has volume snapshots.

To split a mirror component from the mirror:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Select a mirrored volume that is part of the mirror from which you want to split a component.

The Mirrored Volume Details page is displayed.

3. In the Mirror section of the page, select the mirror component that you want to split from the mirror.

The Mirror Component Details page is displayed.

4. Click Split.

A confirmation message and the Mirror Component Details page are displayed.

The Mirrored Volume Details page updates the Mirror section with the Condition of OK, Split Component for the component you selected.

The Volume Summary page now lists the split component as a separate volume with a Condition of Split Component.

Rejoining a Split Mirror Component

You can perform a rejoin operation on a split mirror component and bring it back into the mirror. When you do so, the system immediately begins to resilver (synchronize) the split mirror component to be identical to the other mirror components in the mirror.

To rejoin a split mirror component to the mirror:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Select a mirrored volume.

The Mirrored Volume Details page is displayed.

3. In the Mirror section of the page, click the radio button to the left of the split component that you want to rejoin. Its condition status will be OK, Split Volume.

4. Click Rejoin.

A confirmation message and the Mirrored Volume Details page are displayed.

During the rejoin process, the status of the component is listed as Resilvering.

When the resilvering process is complete, the Mirror section is updated to show that the component is 100% resilvered and that it has a condition of OK.

Reverse Rejoining a Split Mirror Component

You can perform a reverse rejoin operation on a split mirror component, when you want to use its data as the basis for the mirror as a whole. The result is that every mirror component in the mirror becomes identical to the previously split mirror component and the data set represented in the mirror matches that of the previously split mirror component.

To reverse rejoin a split mirror component to the mirror:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select a mirrored volume that is part of the mirror to which you want to reverse rejoin a split component.
The Mirrored Volume Details page is displayed.
3. In the Mirror section of the page, click the radio button to the left of the split component that you want to reverse rejoin.
4. Click Reverse Rejoin.

A confirmation message and the Mirrored Volume Details page are displayed. During the reverse rejoin process, the status of the component is listed as Resilvering. When the resilvering process is complete, the Mirror section is updated to show that the mirror component is 100% resilvered and that it has a condition of OK.

Breaking a Mirror Component

You can remove a mirror component from the mirror and void its association with the other mirror components. The broken mirror component becomes a standalone volume in the system.

You might break a mirror component in order to create an independent copy of an existing volume while the mirror is active and save it as a picture of your data at a specific time.

Note: If you break the last remaining component of a mirror, the result will be a non-mirrored volume with the same name as the mirror. The name of the previous mirror component is no longer used.

To break a mirror component from the mirror:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.

2. Select a mirrored volume that is part of the mirror from which you want to break a component.

The Mirrored Volume Details page is displayed.

3. In the Mirror section of the page, select the mirror component that you want to break from the mirror.

The Mirror Component Details page is displayed.

4. Click Break.

A confirmation message and the Volume Details page for the selected mirror component are displayed. The Volume Summary page now lists the broken mirror component as a separate volume.

Forcing the Break of a Mirror Component

If you know that a mirror component's data has been compromised (during a fire or other catastrophe, for instance) and data integrity is lost, you can force the break of a mirror component, regardless of the percentage of resilvering that has occurred, to protect the integrity of the mirror's data. Just as with a break operation, the mirror component for which you implement a force break becomes a standalone volume in the system; however, the condition of its data is degraded, and it is inaccessible after being removed from the mirror during the resilvering process.

To perform a force break operation for a mirror component:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Select a mirrored volume that is part of the mirror containing the component for which you want to force a break.

The Mirrored Volume Details page is displayed.

3. In the Mirror section of the page, select the mirror component that you want to permanently break from the mirror.

The Mirror Component Details page is displayed.

4. Click Force Break.

A confirmation message and the Volume Details page for the selected component are displayed. The Volume Summary page now lists the broken mirror component as a separate, standalone volume in the system.

Configuring Snapshots

This section describes snapshots. It contains the following subsections:

- [“About Snapshots” on page 103](#)
- [“About Snapshot Reserve Space” on page 104](#)
- [“About Snapshot Rollbacks” on page 105](#)
- [“Planning Snapshots and Snapshot Reserve Space” on page 106](#)
- [“Managing Snapshot Reserve Space” on page 108](#)
- [“Managing Snapshots” on page 111](#)

About Snapshots

A snapshot provides an instantaneous copy of volume data at a specific point in time. Snapshots reside in the same storage domain as the volume for which they are taken (called the parent volume). You can work with snapshots as you would any other volume, with the exception that you cannot take a snapshot of a snapshot. Each snapshot can be accessed independently by other applications.

Snapshots can be read/write-mounted and used for backup, application testing, or data mining, without the need for the primary storage to be taken offline. The right-to-use (RTU) license for snapshots is based on the capacity of primary storage that will be copied.

You can use snapshots for the following:

- To reduce downtime for backup operations

Rather than using online data, back up your snapshots to enable critical transactions to continue during the backup process. Mount the snapshot to the backup server, and then back up the snapshot’s data to tape.
- To perform data analysis and test applications with actual, current data

Use the most recent snapshot for data analysis or testing, instead of working with operating data that might interfere with critical transactions.
- To restart applications from the snapshot volume

If an application problem causes questionable data to be written to the parent volume, restart the application with the last known good snapshot until it is convenient to perform a full recovery.

Note: A snapshot of data alone is not a sufficient means of failure recovery, because it is dependent on the parent volume. Continue to use offline backup methods to create full-volume backup copies.

When a host sends a read request to the snapshot, the system determines whether the requested blocks have changed on the parent volume since the time the snapshot was created. One of the following then occurs:

- If the blocks have changed, the read request is satisfied from the data stored in the snapshot reserve space.
- If the blocks have not changed, the read request is satisfied from the parent volume.

When a write operation occurs for a data block of the parent volume in which the data has not changed since the snapshot's creation, the system does the following:

- Copies the (original) data about to be changed to the snapshot reserve space. Snapshots can accept write operations, and this data is also dependent on the snapshot reserve space.
- Writes the new data to the parent volume. Subsequent rewrites of that same data block on the parent volume do not recopy to the snapshot. In this manner, the snapshot records the data on the parent volume at the time the snapshot is taken.
- Adds a record to the snapshot that indicates the new location for the old data.

To see the current snapshots for a volume, go to the Snapshot Summary page, as described in "Displaying Snapshot Information" on page 111.

About Snapshot Reserve Space

The system stores snapshots in an area associated with the parent volume called snapshot reserve space. Snapshot reserve space can be a simple partition or a concatenated group of virtual disks. For each parent volume that will contain snapshots, only one area of snapshot reserve space is needed. When necessary, you can expand snapshot reserve space.

You determine the size of the snapshot reserve space appropriate for your business practice, based on two factors: the number of snapshots that will be taken of the parent volume, and the anticipated percentage of changed data on the parent volume for each snapshot.

It is usually best to allocate the snapshot reserve space when you create the volume, so as to avoid running out of space and needing to expand snapshot reserve space later; however, if you prefer, you can allocate the snapshot reserve space when you take a snapshot of the volume. The snapshot reserve space data is reclaimed as the oldest snapshot is deleted, resnapped, or used for a rollback and thereby released.

When you allocate snapshot reserve space, you are prompted to specify a write activity from 1% to 100% and a snapshot count from 1 to 32. The calculation of maximum snapshot reserve space allows for up to 32 snapshots, even though only eight snapshots can be active at any given time. You can, therefore, have as many as 24 inactive snapshots that you rotate into activity.

In order to understand why you can select up to 32 snapshots to be taken, though only eight snapshots can be active at any time, consider the following scenario:

Suppose you want to take a snapshot of data on the first of a 31-day month and then take daily snapshots for a week. On the eighth day, you will delete the snapshot that is seven days old and take a new snapshot for the new day, while keeping the snapshot from the beginning of the month. By the end of the month, you will have taken 31 snapshots (one per day on a rotating weekly schedule) plus the one for the first day of the month.

Snapshot reserve space is not reclaimed until the oldest snapshot is deleted; since this doesn't happen for a month, space for 32 snapshots is provisioned. By the end of the month, there are 23 inactive snapshot spaces between the snapshot taken the first day of the month and the most recent seven. This space is reclaimed when the oldest snapshot is deleted and a new monthly snapshot is taken.

Note: As you monitor your storage, use the threshold notifications as a measure of how quickly the snapshot reserve space is being consumed. At several thresholds, the system generates messages to indicate the level of space remaining. The system reports the snapshot full percentage for each volume in 5% increments. When a volume's filled space reaches 70%, the system generates a Warning notification. At the 90%, 95% and 100% thresholds, the system generates a Critical notification.

Once snapshot reserve space is allocated, you can expand it at any time, as described in "Expanding Snapshot Reserve Space" on page 109.

About Snapshot Rollbacks

When you roll back snapshot data, you copy data from a snapshot to the parent volume for which it was taken, resetting the volume's data to become identical to that of the snapshot. This feature is particularly helpful if there are issues with a volume's data.

The rollback snapshot reserve space data can be reclaimed as it becomes unused by other snapshot images when the rollback completes. During the rollback process, additional snapshot reserve space may be consumed to keep the older images up to date.

Note: You cannot roll back a snapshot of a mirrored or replicated volume.

You start the rollback process by selecting the snapshot that you want to roll back; this need not be the newest snapshot.

A background task manages rollbacks. During a rollback, you cannot perform read/write activity on the parent volume or snapshots associated with it. In addition, a rollback operation may take a significant amount of time to complete, depending on the amount of data that is being copied from the snapshot back to the parent volume. When the rollback operation is complete, the create time for the rollback snapshot volume is updated to reflect the time the rollback operation was started.

The snapshot rollback feature allows for the successful completion of scheduled jobs (such as cron jobs), even when a rollback is in progress. You can also create snapshots for volumes other than the parent volume being rolled back.

It is a good idea to make a backup copy of data before issuing a rollback. During a rollback operation, the parent volume data is in flux – it is a composite of the parent volume and the rollback snapshot’s data. The parent volume is dependent upon the snapshot reserve space, data in the rollback snapshot, the metadata, and the snapshot log. Failure of any of these can lead to loss of the original data for the parent volume and loss of the snapshot data being rolled back.

It is important to select the appropriate snapshot policy when performing a rollback. For more information, see “Selecting the Snapshot Policy” on page 111.

Unless you are planning to delete the parent volume and snapshots associated with it, you should not cancel a rollback operation that is in progress. A canceled rollback operation causes the data in the volume and snapshot reserve space to be inconsistent and the condition of the canceled rollback snapshot to be “Non Recoverable Error.”

Planning Snapshots and Snapshot Reserve Space

The first step in planning snapshots is planning the reserve space upon which they depend. You can allocate the snapshot reserve space for a volume when you create the volume, or you can allocate the snapshot reserve space later. A volume with which snapshots are associated is called a parent volume.

Allocating snapshot reserve space when you create the volume helps ensure that there will be sufficient free space for the snapshots. When you allocate snapshot reserve space later, the required amount of free space may not be available; in this case, you will need to delete snapshots, use the resnap operation to retake snapshots, or expand the snapshot reserve space.

Note: Once you have allocated the size of the snapshot reserve space, you cannot decrease it without deleting and then re-creating it.

You must consider two important factors in determining the size of snapshot reserve space:

- The predicted quantity of snapshots.
- The frequency of write operations expected for the parent volume. Only overwritten data is copied into the snapshot reserve space.

For example, assume a 12-Tbyte production database generates 100 Gbytes of changes on a daily basis. If the requirement is to keep daily copies online for a week (seven snapshot copies of the data), the set of copies will require approximately 700 Gbytes of snapshot reserve space.

The Snapshot wizard guides you through the process of allocating snapshot reserve space. The system automatically provisions space for the snapshot metadata, which requires a minimal amount of the size of the parent volume.

Snapshot reserve space is consumed when a block is written to the parent volume that has not been written since the snapshot was created, and when a block is written to the volume for the first time. Estimating the rate at which these activities occur can be difficult. If you are unable to estimate the amount of snapshot reserve space needed, it is recommended that you set the level of activity to 100%, thereby reducing the risk of running out of space.

To plan snapshot reserve space, you need to know:

- Length of snapshot activity

Once a snapshot is taken, subsequent changes to the parent volume data are saved in the snapshot reserve space. Over time, as more blocks of the parent volume are updated, more data is copied to the snapshot reserve space.

- Number of snapshots to be taken

All snapshots of a parent volume share space from the same snapshot reserve space. The more snapshots are taken, the more snapshot reserve space may be required. When you configure snapshot reserve space, specify the size as a percentage of the parent volume.

Note: Keep in mind that you can allocate space for up to 32 snapshots; however, only eight snapshots can be active at any given time.

- Degree of write activity anticipated on the parent volume

Write activity is the number of storage blocks that may change on the parent volume after each snapshot is taken and is a measure of the net change in the parent volume over time. Changes to the parent volume and associated snapshots are stored in the snapshot reserve space.

- Age of the snapshots

The older a snapshot is, the more data is likely to have changed, although the rate of change has probably slowed over time. A snapshot will never consume more space than the size of the parent volume (when 100% of the volume's data has changed since the snapshot was first taken).

If the parent volume is expected to change substantially between snapshots, mirrored volume technologies may be more appropriate, as described in "Configuring Mirrored Volumes" on page 93.

You can resnap a snapshot and replace the old with the new, as described in "Resnapping a Snapshot" on page 113.

You can expand the snapshot reserve space, as described in "Expanding Snapshot Reserve Space" on page 109.

Managing Snapshot Reserve Space

This section describes ways in which you can manage snapshot reserve space. It contains the following subsections:

- ["Allocating Snapshot Reserve Space" on page 108](#)
- ["Expanding Snapshot Reserve Space" on page 109](#)
- ["Deleting Snapshot Reserve Space" on page 110](#)

Allocating Snapshot Reserve Space

You can allocate snapshot reserve space when you create a volume, or after creating a volume. Once you have allocated the size of the snapshot reserve space, you cannot decrease it without deleting and then re-creating it.

You must consider a number of factors and make a number of decisions before allocating snapshot reserve space. For information on planning this allocation, see "Planning Snapshots and Snapshot Reserve Space" on page 106.

To allocate snapshot reserve space after creating a volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Select the volume for which you want to allocate snapshot reserve space.
The Volume Details page is displayed.

3. Click Snapshot Reserve.
4. Specify the settings you want for the following fields:
 - Storage Pool
 - Space Allocated For
 - Expected Write Activity
 - Low – (default) Up to 10% snapshot reserve space needed.
 - Very Little – Up to 25% snapshot reserve space allocated, for up to three active snapshots of average size at any given time.
 - Little – Up to 40% snapshot activity is expected at any given time.
 - Average – Called “Average,” this level of activity may be high for most customers. If you expect snapshot activity to cause a 50% load on snapshot reserve space at any given time, select this level.
 - High – If you plan to use 75% of snapshot reserve space (for up to six or seven active snapshots), select this level.
 - Full – It is unlikely you will use 100% of your snapshot reserve space, for if you rotate snapshots in a cycle, there will be no point at which 100% capacity is reached. However, if you plan to have as many as eight active snapshots at one time, select the Full level of Expected Write Activity.
 - Snapshot Policy
 - Favor parent volume over snapshots (default)

Select this policy for applications like data mining, to avoid problems with the snapshot reserve space that might interfere with business operations on the original data set. Running out of snapshot reserve space will impede the data mining application; however, normal operations will continue.
 - Favor snapshot validity over volume

If it is imperative that you preserve the snapshot data, this policy will protect the integrity of the snapshots and ensure that they are maintained in the system.
5. Click Save.

Expanding Snapshot Reserve Space

You can expand snapshot reserve space in order to make room for more snapshots.

Note: You can expand the snapshot reserve space of a mirrored volume as you would that of any other volume.

To expand snapshot reserve space:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the parent volume for which you want to expand snapshot reserve space.
The Volume Details page is displayed.
3. Click Snapshot Reserve.
4. Increase the values of the settings, as appropriate, for the following fields:
 - Space Allocated For
 - Expected Write Activity
5. Click Save.

Deleting Snapshot Reserve Space

Before you can delete snapshot reserve space, you must delete any snapshots that reside in that reserve space. For instructions, see “Deleting Snapshots” on page 113.

When you have deleted all snapshots for the volume, you can delete the snapshot reserve space.

To delete snapshot reserve space:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the parent volume for which you want to delete snapshot reserve space.
The Volume Details page is displayed.
3. Click Snapshot Reserve.
4. Click Delete Snapshot Reserve.
5. Click Yes to confirm the deletion.
6. Click Save.

Managing Snapshots

This section describes ways in which you can manage snapshots. It contains the following subsections:

- [“Displaying Snapshot Information” on page 111](#)
- [“Selecting the Snapshot Policy” on page 111](#)
- [“Creating a Snapshot” on page 112](#)
- [“Resnapping a Snapshot” on page 113](#)
- [“Deleting Snapshots” on page 113](#)
- [“Rolling Back a Snapshot” on page 114](#)
- [“Canceling a Rollback Job” on page 114](#)

Displaying Snapshot Information

You can display summary and detail information for existing snapshots.

To display snapshot information:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select a snapshot to see its detailed information.
The Snapshot Details page is displayed.

Selecting the Snapshot Policy

Based on your data usage, you can set the snapshot policy to favor the parent volume’s data set or to favor snapshots you take of the data set.

To select the snapshot policy for a volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the parent volume for which you want to select the snapshot policy.
The Volume Details page is displayed.

3. In the Snapshot Reserve Space section of the page, select the snapshot policy you want:

- Favor parent volume over snapshots (default)

Select this policy for applications like data mining, to avoid problems with the snapshot reserve space that might interfere with business operations on the original data set. Running out of snapshot reserve space will impede the data mining application; however, normal operations will continue. If the system runs out of snapshot reserve space, it will continue to accept write commands to the parent volume, but all existing snapshots will become invalid, and all system activity pertaining to these snapshot images will fail.

Note: If this policy is selected for a rollback operation and an error is encountered, the system will attempt to preserve the volume's data until the rollback process is complete.

- Favor snapshot validity over volume

If it is imperative that you preserve the snapshot data, this policy will protect the integrity of the snapshots and ensure that they are maintained in the system. Since problems updating the snapshot reserve space cause all attempted activity performed on the parent volume and snapshots to cease until you can address the problem, this preserves the integrity of the backups.

Note: If this policy is selected for a rollback operation and an error is encountered, the rollback process will be impeded.

4. Click Save.

Creating a Snapshot

You must consider a number of factors and make a number of decisions before creating a snapshot. For more information, see "Planning Snapshots and Snapshot Reserve Space" on page 106.

To create a snapshot:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Volumes.

The Volume Summary page is displayed.

2. Select the check box to the left of the volume for which you want to create a snapshot.

3. Click Snapshot.

The New Snapshot wizard is displayed.

4. Follow the directions in the wizard.

Click Help for more information.

Upon completion of the new snapshot, a confirmation message is displayed.

Note: When you create a snapshot from the Volume Summary page, the system automatically names it (such as “snap050301164439”). To change its name: Select the volume for which you created the snapshot. On the Volume Details page for that volume, go to Additional Information, and click Snapshots. Then, select the snapshot and change the name.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Resnapping a Snapshot

You can replace a previous snapshot with a new one by resnapping the data.

To resnap a snapshot:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the snapshot that you want to resnap.
The Snapshot Details page is displayed.
3. Click Resnap.

Upon completion of the new snapshot, a confirmation message is displayed.

Note: This operation enables you to preserve the World Wide Name (WWN) of the snapshot being replaced.

Deleting Snapshots

Occasionally, you will need to delete snapshots in order to free up snapshot reserve space. In addition, since you can have only eight active snapshots at any given time, you might delete a snapshot to replace it with a new one.

If you are replacing a deleted snapshot right away, you do not need to delete the snapshot you are replacing; you can perform both actions at once by simply resnapping the original snapshot, as described in “Resnapping a Snapshot” on page 113.

To delete a snapshot:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the check box to the left of the snapshot that you want to delete.
This enables the Delete button.
3. Click Delete.
Upon completion of the deletion, a confirmation message is displayed.

Rolling Back a Snapshot

When you roll back snapshot data, the data for the original (parent) volume is reset to be identical to the snapshot.

Rolling back a snapshot for a large parent volume may take a significant amount of time. Cancelling the operation will cause a Non Recoverable error, as described in “About Snapshot Rollbacks” on page 105.

Note: You cannot roll back a snapshot of a mirrored or replicated volume.

To initiate the rollback process:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
2. Select the snapshot that you want to roll back.
The Snapshot Details page is displayed.
3. Click Rollback.
4. When prompted, click Yes to confirm that you want to roll back the data.
A rollback job is created, and the Snapshot Summary page is updated to show that the snapshot’s condition has changed to Snapshot Rollback.

Canceling a Rollback Job

Unless you are planning to delete the parent volume and snapshots associated with it, you should not cancel a rollback operation that is in progress, as described in “About Snapshot Rollbacks” on page 105.

To cancel a rollback job:

1. Click Sun StorEdge 6920 Configuration Service > Jobs > Current.
The Current Job Summary page is displayed.
2. Select the rollback job ID for the job you want to cancel.
The Current Job Details page is displayed.
3. Click the Cancel Job button.
4. Click OK to confirm the cancellation.

Using Data Replication

This section describes data replication. It contains the following subsections:

- [“About Data Replication” on page 116](#)
 - [“Planning for Data Replication” on page 134](#)
 - [“Managing Data Replication” on page 139](#)
-

About Data Replication

This section describes data replication concepts. It contains the following subsections:

- [“About the Replication Process” on page 116](#)
- [“About Replication Sets” on page 118](#)
- [“About Consistency Groups” on page 121](#)
- [“About Replication Links” on page 123](#)
- [“About Replication Set Properties” on page 126](#)
- [“About Primary and Secondary Role Reversal” on page 128](#)
- [“About Replication Modes” on page 129](#)
- [“Reference: Synchronous Versus Asynchronous Replication Modes” on page 131](#)
- [“About Data Replication Operations” on page 132](#)

About the Replication Process

The data replication software is a volume-level replication tool that protects your data. You can use this software to replicate disk volumes between physically separate primary and secondary Sun StorEdge 6920 systems in real time. The software is active while your applications access the data volumes, and it continuously replicates the data to the remote site.

As part of a disaster recovery and business continuance plan, the software enables you to keep up-to-date copies of critical applications at remote sites. You can also rehearse your data recovery strategy to fail applications over to remote sites. Later, you can update the remote site with any changes that occurred on the production data set during the rehearsal, as well as restore any data that was changed on the remote site as part of the rehearsal.

The software replicates data from a primary volume to a secondary volume. The association between the primary and secondary volumes, and a corresponding replication bitmap at each site, make up a replication set. After the volumes in a replication set have been initially synchronized, the software ensures that the primary and secondary volumes contain the same data on an ongoing basis.

Note: There are exceptions if you use asynchronous mode replication. See “About Replication Modes” on page 129 for more information.

Note: Third-party applications can continue to access the primary volume while it is replicating, but not the secondary volume.

When replicating data, the software preserves write order consistency. That is, the software ensures that write operations to the secondary volume occur in the same order as the write operations to the primary volume. This ensures that the data on the secondary volume is consistent with data on the primary volume and does not compromise an attempt to recover the data if a disaster occurs at the primary peer.

If you need to ensure write order consistency across multiple volumes, such as for an application that builds its database on multiple volumes, you can place multiple replication sets into a consistency group. A consistency group enables you to manage several replication sets as one. By using a consistency group, the software maintains write ordering for volumes in a group to ensure that the data on all secondary volumes is a consistent copy of the corresponding primary volumes.

The software transports data between the two Sun StorEdge 6920 systems by means of synchronous replication or asynchronous replication, using either a Fibre Channel (FC) connection or a Gigabit Ethernet network link (replication link). The following illustration shows a data replication over a network link.




FIGURE 4-1 Data Replication Over a Network Link

Note: The system does not provide built-in authentication or encryption for data traveling outside of your data center over a long-distance replication link. It is assumed that customers implementing remote replication strategies using multiple Sun StorEdge 6920 systems will replicate data over secure leased lines or use edge devices to provide encryption and authentication. For help setting up appropriate security, contact Sun Professional Services.

If there is a break in the network or if the secondary peer is unavailable, the software automatically switches to suspended mode, in which it ceases replication and tracks changes to the primary peer in the replication bitmap. When communication is re-established, the software uses the replication bitmap to resynchronize the volumes and returns to replicating the data.

The software can also restore data from a secondary volume to a primary volume by reversing the roles of the primary and secondary peers. Role reversal is a failover technique in which a primary peer failure causes the secondary peer to assume the role of the primary peer. The application software accesses the secondary volume directly until you can correct the failure at the primary peer.

About Replication Sets

A replication set includes:

- A volume residing on a Sun StorEdge 6920 system and a reference to a volume residing on another, physically separate Sun StorEdge 6920 system. One system is the primary peer, which copies the data, and the other system is the secondary peer, which is the recipient of the data.
- A replication bitmap for logging purposes.
- The communication mode between both systems.
- The role that the peer plays within the replication set, either as a primary or as a secondary peer.

The following figure shows the relationship between the two peers and their corresponding replication sets.




FIGURE 4-2 Relationship Between Two Peers and Their Corresponding Replication Sets

The system administrator at each site must create and configure a replication set on each system. The replication set definition for the secondary peer must be equivalent to that of the replication set for the primary peer.

You can update the secondary volumes synchronously in real time or asynchronously using a store-and-forward technique. Typically, a primary volume is first explicitly copied to a designated secondary volume to establish matching contents. As applications write to the primary volume, the data replication software copies the changes from the primary volume to the secondary volume, keeping the two images consistent.

The replication set also includes the following:

- A replication bitmap volume on each system.

The replication bitmap tracks write operations and differences between the volumes. The primary volume's bitmap records write actions issued at the primary peer. The secondary peer's replication set also includes a replication bitmap in case you initiate a role reversal and the secondary peer becomes the primary peer. The replication bitmap defines the differences between the primary and secondary peers. This enables the software to resynchronize only the blocks that have changed since the last synchronization.

- If you choose asynchronous mode replication, an asynchronous queue associated with each peer.

After you create a replication set, you can perform volume operations, such as:

- Extend replicated volumes, by adding storage first to the secondary volume and then to the primary volume. The system increases the size of the replication bitmaps accordingly.

Note: If the primary or secondary volume in a replication set is a legacy volume, you cannot extend the volume. If the secondary volume in a replication set is a legacy volume, you cannot extend the primary volume beyond the size of the legacy volume.

- Create snapshots.

You can create a snapshot on any volume within a replication set at any time.

You can also enable options so that the software creates volume snapshots on both systems before synchronizing the volumes, or so that the software creates a volume snapshot of the secondary volume whenever it creates a volume snapshot of the primary volume.

- Replicate user data to and from a legacy volume.
- Delete a replication set.

When you delete a replication set, you remove the association between the primary and secondary volumes, and the volumes revert to independent volumes.

Note: You cannot mirror a volume that is also in a replication set.

If the primary volume becomes unavailable, the secondary volume can assume the role of primary volume. This role reversal allows applications to continue their operations by using the newly designated primary volume. When the former primary volume is again available, you must synchronize it with the more recent data on the other volume to restore the functions of the replication set pair.

Before you create a replication set, consider the restrictions and other factors described in the following table.

TABLE 4-1 Replication Set Restrictions and Considerations

Regarding	Note
Ports	<p>You must enable the ports on both the local and remote Sun StorEdge 6920 systems for data replication before you run the Create Replication Set wizard. If the port is a Gigabit Ethernet port, you must also create a replication link. For more information about enabling ports for replication, see “Configuring System Ports” on page 175.</p> <p>Note: The replication peer port must be initialized before the replication set can initialize. Be sure you have entered the Replication Peer WWN (16 hexadecimal digits) before creating a replication set.</p>

TABLE 4-1 Replication Set Restrictions and Considerations (*Continued*)

Regarding	Note
Storage domains	<p>All elements of a replication set (the local volume, the replication bitmap, and the optional asynchronous queue) must be from the same storage domain.</p> <p>Note: A replication set is automatically associated with the storage domain of the local volume.</p>
Volumes	<ul style="list-style-type: none">• You can associate replication sets with configured volumes only.• Volumes must be in storage domains that are named the same on each system. Note: Though the storage domains on both systems have the same name, they are not the same storage domain.• You can use simple or legacy volumes to create a replication set. Note: You cannot use a mirrored volume to create a replication set.• You cannot use a snapshot to create a replication set, but you can use a volume that includes snapshots.• You can use only unmapped volumes as secondary volumes.• A volume can be the secondary volume for only one replication set.• The size of the secondary volume must be equal to or greater than the size of the corresponding primary volume. If you initiate resynchronization on a replication set in which the secondary volume is smaller than the primary volume, the operation will fail.
Replication sets	<ul style="list-style-type: none">• A replication set can be configured to communicate with only one peer system.• A replication set can be a member of only one consistency group at a time.• A replication set must be configured with the same role as the consistency group to which you want to add it.• The replication set specifications at the primary and secondary peers must match and complement each other.• A system that includes two storage resource card (SRC) sets supports up to 128 replication sets. A system that includes four SRC sets supports up to 256 replication sets. <p>Note: Multi-hop and one-to-many data replication are not supported.</p>

About Consistency Groups

To ensure write order consistency across multiple volumes, you can group multiple replication sets into a consistency group. A consistency group is a collection of replication sets that have the same group name, primary and secondary roles, and replication mode. Mixed groups, in which replication modes are asynchronous for one replication set and synchronous for another replication set, are not allowed.

Note: As with replication sets, the system administrator at each site must each create and configure a consistency group for the respective peers. The consistency group for the secondary peer must complement the consistency group for the primary peer.

When you perform an operation on a consistency group, the operation applies to all the replication sets, and consequently their volumes, in the consistency group. If you make a change to a consistency group, the change occurs on every replication set in a consistency group; if an operation fails on a single replication set in the consistency group, it fails on every replication set in the consistency group.

Note: Volume snapshot operations are the exception. You must create a snapshot of each replication set individually.

When you configure a consistency group, the system preserves write ordering among the volumes in the replication sets. Because you control the replication sets as a single unit, data replication operations are executed on every member of the consistency group. Write operations to the secondary volume occur in the same order as the write operations to the primary volume. By using a consistency group, the software maintains write ordering among volumes in a group to ensure that the data on each secondary volume is a consistent copy of the corresponding primary volume.

Before you create a consistency group or place a replication set in a consistency group, consider the restrictions and other factors in the following table.

TABLE 4-2 Consistency Group Restrictions and Considerations

Regarding	Note
Write order	If you need to ensure write order consistency for a set of volumes, you must place replication sets in a consistency group. You can create the consistency group at the same time that you create the first replication set or you can add an existing replication set to the consistency group from the replication set's Details page.
Volumes	<ul style="list-style-type: none">• All primary volumes in a consistency group must reside on the same primary peer. You cannot preserve write ordering when you have primary volumes originating on different systems.• All secondary volumes in a consistency group must reside on the same secondary peer. You cannot place secondary volumes on different systems and expect to preserve write ordering among them.

TABLE 4-2 Consistency Group Restrictions and Considerations (*Continued*)

Regarding	Note
Replication sets	<ul style="list-style-type: none">• Replication sets in a consistency group must be part of the same storage domain.• You must define identical consistency groups for both the primary and secondary peers in order to preserve write ordering on data arriving at the secondary peer.• Replication sets in a consistency group must have the same replication mode.• Replication sets in a consistency group must have the same role. If you attempt to add a replication set to a consistency group with a different role, the operation will fail.• If you add a replication set that is configured for asynchronous replication to a consistency group, the local asynchronous queue is deleted. The consistency group has its own asynchronous queue.• Each consistency group supports up to 128 replication sets.

A best practice is to plan for consistency groups ahead of time and create them at the same time that you create replication sets. You cannot directly combine replication sets in a consistency group.

If you plan to create a new replication set and you need to group it with an existing replication set in a new consistency group, follow these steps:

1. Use the Create Replication Set wizard to create the new replication set and a new consistency group.
2. From the existing replication set's Details page, move the replication set into the newly created consistency group to join the new replication set.

About Replication Links

A replication link is a logical and physical connection between two Sun StorEdge 6920 systems that allows for data replication. A replication link transports data between the primary and secondary peers. This link transfers data as well as replication control commands.

You can use both Fibre Channel (FC) and Gigabit Ethernet ports for data replication. You must enable the same types of ports on both systems to establish the replication link.

Note: You can configure only two replication links at a time, and the replication links must both be on either FC ports or Gigabit Ethernet ports. You cannot mix port types.

If you use FC ports for data replication, you need not perform any other configuration tasks. However, you must configure any FC switches that you use to make the connection to the remote site for long distance operations and apply zoning practices (see the following figure).

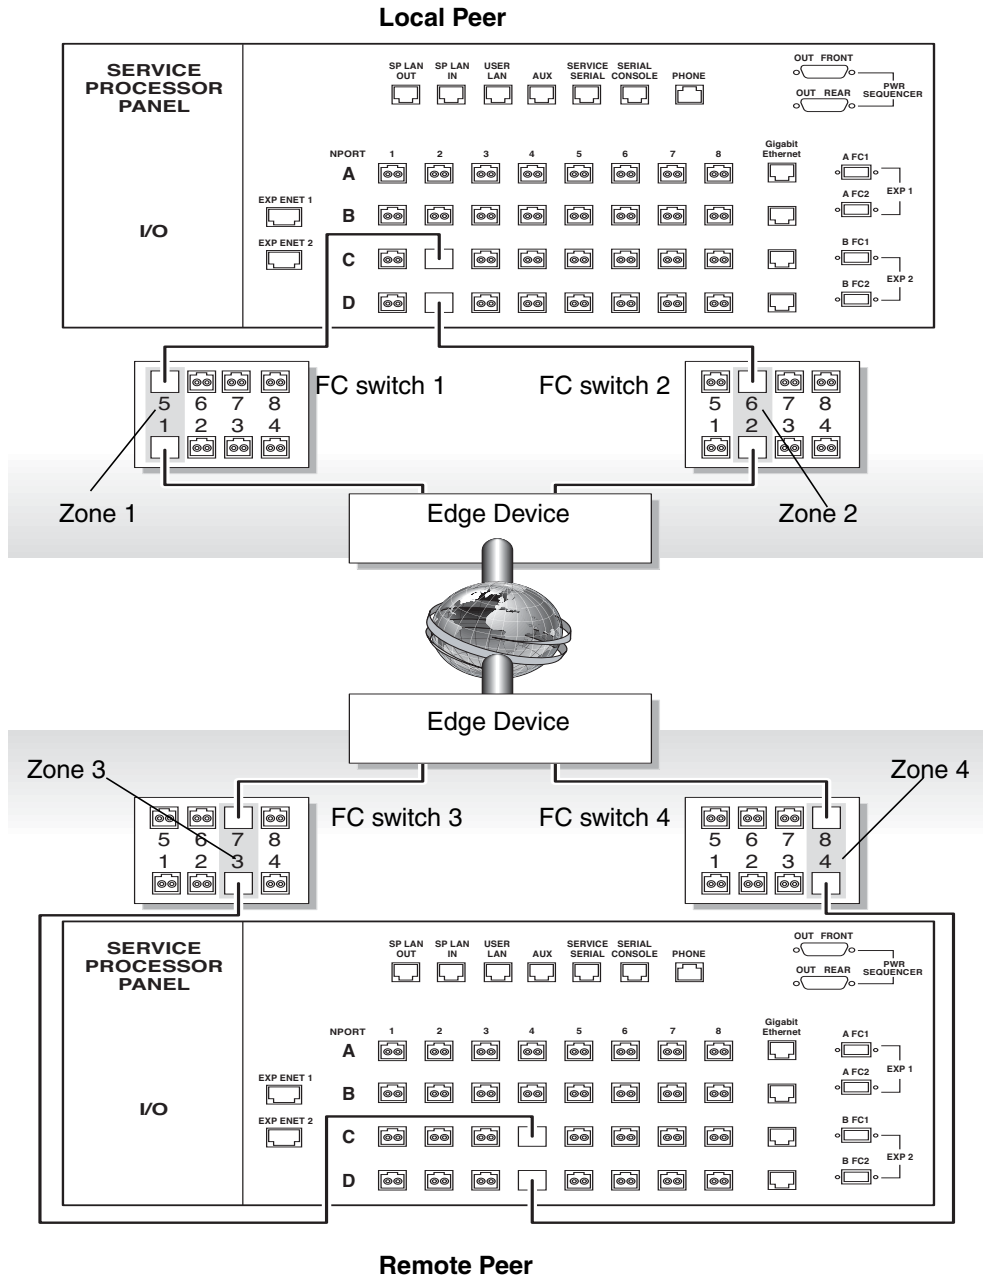


FIGURE 4-3 Example FC Switch Configuration

Note: The system does not provide built-in authentication/encryption for data traveling outside of your data center over a long-distance replication link. It is assumed that customers implementing remote replication strategies using multiple Sun StorEdge 6920 systems will replicate data over secure leased lines or use edge devices to provide encryption and authentication. For assistance with setting up appropriate security, contact Sun Professional Services.

See the FC switch vendor's documentation for information about operating over long distances.

If you use Gigabit Ethernet ports, you must explicitly configure the ports to create the replication link. Configure a replication link by running the Enable Gigabit Ethernet Port Replication wizard from the Gigabit Ethernet Port Details page. For more information, see "Enabling Ports for Data Replication" on page 177.

About Replication Set Properties

When you define a replication set or a consistency group, you set replication properties. To set these properties, do the following:

- Set the role of the replication set on one system to primary and the role of the complementary replication set on the other system to secondary.
- Set the replication mode to synchronous or asynchronous. See "About Replication Modes" on page 129 for more information.
- Set the priority of the synchronization rate relative to I/O activity to high, medium, or low.
- If appropriate for your environment, enable the autosynchronization option.

Autosynchronization is an alternative to manual synchronization. The autosynchronize option supports both replication sets and consistency groups. If you enable the autosynchronize option on the primary peer, the software synchronizes the volumes on both peers and resumes replication as soon as possible. For example, if a network link fails and causes the software to cease replication, synchronization will occur when the link is re-established.

If autosynchronization is enabled, the software attempts to synchronize the secondary volume with the primary volumes if the replication set or consistency group is placed in suspended mode because of a link failure or system shutdown. The software does not perform an autosynchronization operation if the replication set or consistency group was placed in suspended mode through either of the following methods:

- You manually set the replication set or consistency group to suspended mode.
- The asynchronous queue exceeded its limits while the replication link was active.

Note: When a replication set or consistency group resumes data replication, the autosynchronize option is not disabled.

Follow these instructions with respect to the autosynchronize option:

- If you choose the Fast Start option, disable the autosynchronize option so that you can start the synchronization operation manually when you have restored the secondary volume and it is ready for replication.
- Keep the secondary peer as consistent as possible to avoid having the primary and the secondary peers inconsistent with each other, and difficult to recover. For this reason, autosynchronization is disabled by default.
- If appropriate for your environment, enable the option to create volume snapshots automatically.

Disaster can occur during any phase of data replication, such as during a synchronization operation. Although data replication does not affect the integrity of the primary data, the secondary data is vulnerable during synchronization because write order is not preserved. Therefore, to ensure a high level of data integrity on both peers during normal operations or data recovery, create volume snapshots before you perform a synchronization operation so that you always have a consistent copy of your data. If a failure occurs, the snapshot provides a point-in-time data copy from which your data can be restored.

The software provides two options for creating volume snapshots automatically:

- When you choose the Snap Before Synchronization option, the software creates a volume snapshot at the secondary peer prior to the synchronization operation. When the synchronization operation is completed successfully, the volume snapshot is deleted. If the synchronization operation is interrupted, you can use the volume snapshot to roll back the secondary volume to a consistent state.
- When you choose the Snap With Primary option, the software creates two volume snapshots, one of the primary volume and one of the secondary volume. The resulting snapshots at the primary and secondary peers represent the same point in time in your data.

If you are using an asynchronous queue, the snapshot of the secondary volume occurs after data in the queue has been written to the secondary volume, and before any writes have been added to the queue since the snapshot operation. This ensures that the snapshots at the primary and secondary peers represent the same point in time in the data.

When you enable the volume snapshots options, consider the following:

- Volume snapshots guarantee a known, write-ordered copy of your user data in the case of a network link failure during a synchronization operation, as long as the snapshot operation is successful. Make sure that you have allocated sufficient snapshot reserve space so that snapshot operations succeed.
- There is some degradation in performance if you choose the Snap Before Synchronization option.

- If you create many volume snapshots at the primary peer and do not want to consume resources at the secondary peer, disable the Snap With Primary option.
- Because you cannot access the secondary volume for read and write operations unless it is in suspended mode, you can create a volume snapshot of the secondary volume, which you can use to access the user data while the volume is being replicated.
- If you enable the Snap Before Synchronization option for a secondary replication set, and the secondary volume does not include snapshot reserve space, the software returns the following error message to the primary peer when you try to resume data replication:

PIT failed for sync startup

Add snapshot reserve space to the secondary volume to resume data replication successfully. If you disable the Snap Before Synchronization option, the secondary volume will be only partially replicated on the secondary replication set.

About Primary and Secondary Role Reversal

In case of disaster or link failure, you can swap roles to continue to provide access to your critical data. This is a useful failover technique that enables the volumes in a replication set to continue business-critical operations when the primary peer fails.

A role reversal causes the primary peer to become the secondary peer and the secondary peer to become the primary peer. The application software can run on what was originally the secondary peer. When the original primary peer failure has been rectified, another role reversal returns both peers to their original roles.

A best practice is to create volume snapshots of the secondary volume regularly in case its data becomes inconsistent due to an in-progress or terminated synchronization operation. If you have a volume snapshot available, you can perform a rollback operation from the snapshot to restore the data.

Note: To perform a rollback operation, you must first delete the replication set on the secondary peer. After you complete the rollback operation, create the replication set again.

Caution: If I/O operations to the primary and secondary peers occur during a role reversal, data is lost when you initiate a synchronization operation. That is, the current secondary volume is brought into synchronization with the current primary volume, and any writes that have been written to the secondary volume are lost.

About Replication Modes

The replication mode is a user-selectable property that defines the communication mode for a replication set. The software supports two modes of data replication:

- Synchronous mode

In synchronous mode replication, a write operation to the primary volume is not confirmed as complete until the remote volume has been updated. Synchronous replication forces the software to wait until the primary peer receives an acknowledgment of the receipt of the data from the secondary volume before returning to the application.

- Asynchronous mode

In asynchronous mode replication, data is written to the primary volume and to a local asynchronous queue. A write operation is confirmed as complete before the remote volume has been updated. Later, write operations that have accumulated in the asynchronous queue are forwarded in sequence to the remote peer. Asynchronous replication enables the data replication software to return to the peer as soon as the write operation has been completed on the primary volume and has been placed on a per-volume queue for the secondary peer. The secondary peer receives the queued requests in the order in which they were written. After the write operation has been completed at the secondary peer, notification is sent to the primary peer.

The asynchronous queue exists to absorb bursts of application writes. You can select how the asynchronous queue operates when it becomes full and causes application writes to wait for room in the queue:

- Blocking mode – If the asynchronous queue fills, all writes to the primary volume and replication writes to the secondary volume are delayed until the queue drains enough to allow for a write to occur. Blocking mode, which is the default option, ensures write ordering of the data to the secondary peer. If the asynchronous queue fills with the blocking option set, response time to the application might be affected. Write operations to the secondary volume must be acknowledged before being removed from the queue on the primary peer, so they can prevent, or block, further write operations to the queue until space is available.
- Suspended mode – If the asynchronous queue fills, the software discontinues data replication and no longer records writes in the queue. Instead, the software records data block changes in the replication bitmap. The application's writes are not blocked, but write ordering is lost when the software is in suspended mode. However, the application sees no significant degradation in response time.

To resume replication, you must first synchronize the secondary volume with the primary volume. For information about replication operations, see “About Data Replication Operations” on page 132.

Consider the following when you choose asynchronous mode replication:

- All the volumes in a consistency group share a single asynchronous queue.
- If you choose the suspend option, when the queue is full, the software switches to suspended mode. Write ordering is not preserved; however, application write operations are not impacted because of a full queue.
- The minimum size of an asynchronous queue for a consistency group is 16 Mbytes.
- You must choose the proper queue size for your environment. If the asynchronous queue fills, subsequent write operations must wait to be placed in the queue. As a result, the application response time increases. To help improve the response time for the application, increase the asynchronous queue size based on its usage.

If you need to extend the size of the asynchronous queue:

- a. Place the replication set or consistency group into suspended mode.
 - b. Go to the replication set or consistency group Details page and use the pull-down menu to change the asynchronous queue size.
 - c. Initiate a synchronize operation for the replication set or consistency group to synchronize both peers and resume replication.
- You can limit when the queue is considered full by the number of queued disk blocks or length of time an entry is in the queue. To set asynchronous queue parameters, use the Create Replication Set wizard or make changes on the replication set or consistency group Details page.
 - Asynchronous mode accommodates bursts of write activity in which the write rate exceeds the replication link's bandwidth. The asynchronous queue must be sufficient in size to handle bursts of write traffic associated with the application peak write periods. A large queue can handle prolonged bursts of write activity, but this activity causes the secondary peer to become further out of synchronization with the primary peer.
 - If you add a replication set that is configured for asynchronous replication to a consistency group, that replication set's own queue is deleted.
 - Because of the nature of an asynchronous queue, the secondary volume will always be somewhat out of date with the primary volume. How far out of date the secondary volume is compared with the primary volume depends on how much data there is in the asynchronous queue at the time, as well as on the latency of the link.
 - If you choose the block option, and the queue becomes full, writes are blocked until the queue drains. The software maintains write ordering; however, application write operations are impacted.

You can change the replication mode at any time during the life of a replication set. However, you must first place the replication set in suspended mode. If the replication set is a member of a consistency group, you must place the consistency group in suspended mode.

Note: If a replication set is a member of a consistency group, you cannot change the replication mode of the replication set. The replication set's attributes must match those of the consistency group.

Reference: Synchronous Versus Asynchronous Replication Modes

When you choose a replication mode, consider the differences between synchronous and asynchronous replication described in the following table.

TABLE 4-3 Replication Mode Considerations

Characteristic	Synchronous Replication	Asynchronous Replication
Write operations	A write operation to the primary volume is not complete until the user data is written to the secondary peer.	The write operation is complete as soon as the user data is written to the primary volume and to the asynchronous queue.
Network response time	The response time depends on the network. The network latency must be low enough that the application response time is not affected dramatically by the time of the network round trip of each write operation. Also, the bandwidth of the network must be sufficient to handle the amount of write traffic generated during the application's peak write period. If the network cannot handle the write traffic at any time, the application response time will be affected.	The response time does not depend on the network, as long as the queue is not filled. The bandwidth of the network link must be able to handle the write traffic generated during the application's average write period. During the application peak write phase, the excess write operations are written to the local asynchronous queue and then written to the secondary peer at a later time when the network traffic allows. The application response time can be minimized during bursts of write traffic above the network limit as long as the asynchronous queue is properly sized. Note: Response time to the application might be affected if the asynchronous queue fills with the blocking option set.

TABLE 4-3 Replication Mode Considerations (*Continued*)

Characteristic	Synchronous Replication	Asynchronous Replication
Availability	Synchronous replication provides high data availability at the expense of longer latency across the wide area network (WAN).	Asynchronous replication offers higher performance and availability in exchange for delayed synchronization of all copies of data.
Write response time	Synchronous replication might exhibit an increased write response time, especially for large data sets or long-distance replication (where write operations can incur additional latency because of the time required to transfer data and return acknowledgments).	Asynchronous replication can have a longer response time than synchronous replication.
Write order	Because every data change is replicated to the secondary peer in real time, the secondary peer is write order consistent with the primary peer. Write ordering across volumes is preserved at the secondary peer.	The secondary peer lags behind the primary peer, depending on the latency, queue size, and the number of writes. Write ordering is preserved across volumes at the secondary peer when volumes are grouped at the primary peer.
Data loss	Synchronous replication is best for data volumes that have zero tolerance of data loss at the secondary peer.	Asynchronous replication is best for data volumes that have some level of tolerance for data loss at the secondary peer.

About Data Replication Operations

You can choose one of the following data replication operations from either the Replication Set Details Page or the Consistency Group Details Page:

- Resume, which provides the following two options:
 - Normal synchronization – A copy operation of differences between the primary and secondary volumes. Synchronization also enables concurrent replication from the primary volume to the secondary volume so that any new write operations to the primary volume are also replicated to the secondary volume. After synchronization is completed, the volumes are write order consistent with each other.
 - Full synchronization– A complete volume-to-volume copy operation, which is the most time-consuming of the synchronization operations. In most cases, a secondary volume is synchronized from its source primary volume. A full synchronization is typically performed only when the replication set is damaged due to a disaster or has questionable data integrity.

Before you initiate a resume operation, consider the following:

- You cannot initiate a resume operation for a replication set or consistency group with the role of secondary.
- When you initiate a resume operation, it is possible that large quantities of I/O will occur over the replication links as the volumes are brought into synchronization. Make sure that you don't impact bandwidth. For example, avoid creating a new replication set across a limited-bandwidth link.
- A resume operation creates a job, which enables you to track the progress of the synchronization. You can follow the progress of the job and keep a history of its operation from the Job Summary page. The Replication Progress fields on the Replication Set Summary page and each replication set's Details page also report the progress of the synchronization. However, the software updates the job's progress more frequently on the Job Summary page than it does in the Replication Progress fields. Therefore, the Job Summary page might report that the synchronization has progressed more than what the Replication Progress field reports.

Note: If the job is completed quickly, the Job Details page might report that a synchronization job has been started although the page banner does not display it as a running job. The Historical Job Summary page lists the completed job.

- Suspend, which provides the following two options:
 - Normal – An operation to stop replication or synchronization. While in suspended mode, the software updates only the replication bitmap at the primary peer. No replication occurs. At a later time, such as when an operation is completed or a link is re-established, the volumes are synchronized, either through the autosynchronize option or manually. The replication bitmaps at the primary and secondary peers are compared, and the changed blocks for the primary peer volume are copied to the secondary volume. After the volumes are synchronized, replication starts again.
 - Fast Start – An operation that signifies that you are performing the initial synchronization manually by means of a backup tape. You initiate the Fast Start option at the primary peer before you create the backup tape and initiate the Fast Start option at the secondary peer after you restore the backup tape.

Note: When you are ready to initiate a synchronization operation after being in Fast Start mode, you must perform the synchronization operation manually. To ensure that synchronization does not occur at an inappropriate time, the software will not perform a synchronization operation automatically.

When network latencies justify it, you can perform the initial synchronization by backing up a source or primary volume on magnetic tape at one site and then restoring the volume from the tape at the other site. Make sure that the backup copy is a physical copy (for example, one made using the `dd(1M)`

command) and not a logical copy (for example, one made using the `tar(1M)` or `cpio(1M)` commands). The copies must have identical blocks, not just identical files.

Note: If a replication set or consistency group is already in suspended mode, the normal suspend option is not available; you can only select the Fast Start option.

Use the suspend operation as follows:

- Use the Fast Start option of the suspend operation to minimize data replication I/O traffic when you set up a copy of the data at a new site. You can initialize the remote peer by shipping a backup tape that contains a copy of the user data and loading the backup tape directly onto the remote system. This option is useful if you have an extensive amount of user data, are concerned about the cost of the peer link, or are concerned with the speed of the peer link.
- Use a suspend operation to save on telecommunications or connection costs. However, you risk the cost of data loss. If you lose the primary peer, you do not have the data at the secondary peer that was written to the primary peer.

The replication sets can go into suspended mode for either planned or unplanned reasons, such as when you make a change to a consistency group or a network link failure occurs. When the software resumes data replication from suspended mode, the following occurs:

1. While the software is in suspended mode, there might be a physical connection. However, there is no communication and no data replication between the two peers.
2. When synchronization is initiated, there is communication between the two peers, but write order consistency has not been restored.
3. When the peers are synchronized, there is communication between the two peers, and write order consistency is restored. If the replication mode is asynchronous, there might be a time delay, but the data will be write order consistent.

Planning for Data Replication

This section describes planning for data replication. It contains the following subsections:

- [“General Planning Considerations” on page 135](#)
- [“Reference: Data Replication Tasks” on page 136](#)
- [“Planning to Create Replication Sets and Consistency Groups” on page 137](#)

General Planning Considerations

Replicating data and modifying replication options and operations are a significant change to your system's configuration. You must plan accordingly before you perform data replication operations. Consider the following:

- **Business needs** – When you decide to replicate your business data, consider the maximum delay: How long out of date can you allow the data on the secondary peer to become? This determines the replication mode and snapshot scheduling. Additionally, it is very important to know whether the applications that you are replicating require the write operations to the secondary volume to be replicated in the correct order.
- **Application write load** – Understanding the average and peak write loads is critical to determining the type of network connection required between the primary and secondary peers. To make decisions about the configuration, collect the following information:
 - **The average rate and size of data write operations**

The average rate is the number of data write operations while the application is under typical load. Application read operations are not important to the provisioning and planning of your remote replication.
 - **The peak rate and size of data write operations**

The peak rate is the largest amount of data written by the application over a measured duration.
 - **The duration and frequency of the peak write rate**

The duration is how long the peak write rate lasts, and the frequency is how often this condition occurs.
- **Network characteristics** – The most important network properties to consider are the network bandwidth and the network latency between the primary and secondary peers.

Reference: Data Replication Tasks

The following table provides a checklist of data replication tasks to consider:

TABLE 4-4 Checklist of Data Replication Tasks to Consider



 To Do	
	<p>Consider network latency, bandwidth, and security when you choose the path through the storage area network (SAN) or wide-area network (WAN) that you will use for data replication.</p> <p>Note: As the software synchronizes the volumes within a replication set, there might be a large number of additional I/O operations, which might impact performance.</p>
	<p>Make sure that the remote Sun StorEdge 6920 system is available to the your system by means of an existing physical link.</p>
	<p>Log on to the system using the <code>storage</code> user role.</p>
	<p>Configure both the primary and secondary volumes as you would any other volume, and make sure that they have identical configurations. The capacity of the volumes must be precisely the same.</p>
	<p>From the Volume and Port Summary pages, record the:</p> <ul style="list-style-type: none">• WWN of both the local and remote Sun StorEdge 6920 Systems.• WWN of both the local and remote volumes that will make up the replication set• IP addresses of both the local and remote ports, if you are using a Gigabit Ethernet port. <p>Note: If you enter an incorrect WWN in the Create Replication Set wizard, you must delete the replication set and create a new one.</p>
	<p>For configurations that use Fibre Chanel (FC) ports, configure the FC switches that provide the connection between peers for long distance operations. See the FC switch vendor's documentation for information about operating over long distances.</p>
	<p>Enable the FC or Gigabit Ethernet ports you plan to use on both the local and remote Sun StorEdge 6920 systems for replication. If the port is a Gigabit Ethernet port, at both sites, complete the Enable Gigabit Ethernet Replication Port wizard to create a replication link. See "Enabling Ports for Data Replication" on page 177 for information about how to enable Fibre Channel (FC) and Gigabit Ethernet ports for data replication.</p>

TABLE 4-4 Checklist of Data Replication Tasks to Consider (*Continued*)

	 To Do
	<p>Before creating a replication set, consider:</p> <ul style="list-style-type: none"> • Whether your application coordinates its data across several volumes and you plan to use asynchronous replication, and consequently require your replication sets to be in consistency groups. • Replication roles, modes, and priorities. • Asynchronous queue properties. • Volume snapshot requirements. • The replication peer port must be initialized before the replication set can initialize. Be sure you have entered the Replication Peer WWN (16 hexadecimal digits) before creating a replication set.
	<p>Use the Create Replication Set wizard to do one of the following:</p> <ul style="list-style-type: none"> • Create a replication set • Create a replication set and add it to an existing consistency group • Create a consistency group at the same time as you create a replication set <p>Make sure the replication sets and consistency groups for both peers are configured identically.</p>
	<p>Consider the different operations involved with data replication:</p> <ul style="list-style-type: none"> • Normal synchronization operation • Full synchronization operation • Normal suspend operation • Fast Start operation • Role reversal

Planning to Create Replication Sets and Consistency Groups

To replicate data to a secondary peer, you must first run the Create Replication Set wizard to create replication sets and consistency groups. The Create Replication Set wizard enables you to create a consistency group at the same time that you create a replication set.

Before you run the wizard, you must do the following:

- Make sure that you have configured volumes and enabled ports for data replication on both systems. See “Reference: Data Replication Tasks” on page 136 for guidelines to prepare for data replication.

- Coordinate with the system administrator at the remote site to ensure that you both define volumes, replication sets, and consistency groups with properties that complement each other appropriately. For guidelines on creating replication sets and consistency groups, see “About Replication Sets” on page 118 and “About Consistency Groups” on page 121.
- Determine the World Wide Name (WWN) for the remote peer and the secondary volume. You can find these WWNs on the Volume and Port Summary pages for the remote peer. For more information, see “Displaying Volume Information” on page 89 and “Displaying Port Information” on page 176.
- Determine whether you need to preserve write order consistency across volumes or want to manage multiple volumes as a group. If so, create a consistency group at the same time as you create the first replication set; add subsequent replication sets to the existing consistency group. Make sure that the consistency group definition matches that of the replication set you plan to add. For guidelines on creating consistency groups, see “About Consistency Groups” on page 121.
- Define the replication set or consistency group:
 - Assign the primary role to the replication set or consistency group, and ensure that the system administrator at the remote site assigns the corresponding replication set or consistency group with the secondary role. For information about changing roles for existing replication sets and consistency groups, see “About Primary and Secondary Role Reversal” on page 128.
 - Determine whether synchronous or asynchronous data replication mode is the best method of data communication for your application. For more information about replication modes, see “About Replication Modes” on page 129.
 - If you choose asynchronous mode, determine the asynchronous queue parameters and consider how they will affect performance.
 - Determine whether you want the system to assign a higher or lower priority to the rate of synchronization relative to I/O activity. A high synchronization rate might lead to degraded I/O performance, so you can choose a faster synchronization rate if your environment can sustain slower I/O activity. Conversely, you can choose a slower synchronization rate to ensure quicker I/O activity.
 - Determine whether you want the system to synchronize the volumes on both peers whenever there is a working link, or whether you want to initiate synchronization operations manually. For more information about the autosynchronization option, see “About Data Replication Operations” on page 132 and “About Replication Set Properties” on page 126.
 - Determine whether you want to create volume snapshots whenever you synchronize the volumes on both peers. For more information about how the software behaves when you create volume snapshots within a replication set or consistency group, see “About Replication Set Properties” on page 126. You must also ensure that you have allocated sufficient snapshot reserve space. For more information, see “About Snapshot Reserve Space” on page 104.

- If you choose asynchronous mode replication, determine the behavior of the asynchronous queue and the storage pool from which you want to create the queue. For more information, see “About Replication Modes” on page 129.
- Determine the pool from which you want to create the replication bitmap.

Managing Data Replication

This section describes data replication tasks. It contains the following subsections:

- [“Displaying Replication Link Information” on page 139](#)
- [“Displaying Replication Set Information” on page 140](#)
- [“Displaying Consistency Group Information” on page 140](#)
- [“Creating Replication Sets and Consistency Groups” on page 141](#)
- [“Deleting Replication Sets” on page 141](#)
- [“Adding a Replication Set to a Consistency Group” on page 142](#)
- [“Combining Replication Sets in a Consistency Group” on page 143](#)
- [“Removing a Replication Set From a Consistency Group” on page 144](#)
- [“Deleting Consistency Groups” on page 145](#)
- [“Reversing Roles” on page 145](#)
- [“Changing Replication Modes” on page 146](#)
- [“Replicating Data” on page 147](#)
- [“Synchronizing Data Using a Backup Tape” on page 148](#)
- [“Troubleshooting Data Replication” on page 148](#)

Displaying Replication Link Information

You and the system administrator at the remote site can display replication link information to ensure that there are established replication links between peers.

To display information about replication links on Gigabit Ethernet ports on your system:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.
The Port Summary page is displayed.

2. Select the Gigabit Ethernet port for which you want to see replication link information.

The Gigabit Ethernet Port Details page is displayed.

3. In the Additional Information section, click Replication Links.

The Replication Link Summary page is displayed.

Displaying Replication Set Information

You and the system administrator at the remote site can display replication set information to ensure that there is a corresponding replication set established at each peer.

To display information about the replication sets on your system:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Click a replication set.

The Replication Set Details page is displayed.

Note: If the replication set belongs to a consistency group, you can click the consistency group to display the Consistency Group Details page.

Displaying Consistency Group Information

You and the system administrator at the secondary peer can display consistency group information to ensure that there is a corresponding consistency group established at each peer.

To display information about the consistency groups on your system:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. In the Consistency Group column, click a consistency group.

The Consistency Group Details page is displayed.

You can also access the Consistency Group Details page from a replication set's Details page. Click the consistency group name displayed after the Consistency Group label.

Creating Replication Sets and Consistency Groups

You must consider a number of factors and make a number of decisions before creating a replication set or consistency group. For information on planning a replication set or consistency group, see “Planning for Data Replication” on page 134.

The Create Replication Set wizard enables you to:

- Create a replication set that is not in a consistency group
- Create a replication set, and create a consistency group in which the replication set will be placed
- Create a replication set and add it to an existing consistency group

To create a replication set or a consistency group:

1. Click Sun StorEdge Configuration Service > Logical Storage > Volumes.
The Volume Summary page is displayed.
2. Click the name of a volume you want to replicate to the remote peer.
The Volume Details page for the selected volume is displayed.
3. Click Replicate.
The Create Replication Set wizard is displayed.
4. Follow the steps in the wizard.
Click the Help tab in the wizard for more information.

Deleting Replication Sets

When you delete the replication set, the primary volume stops replicating data to the secondary volume. The association between the primary and secondary volumes is removed, and both revert to conventional volumes. The system reclaims the replication bitmap storage and any asynchronous queue that might have existed as free storage. The data stored on the volumes is not affected, and the volumes can remain mapped.

If you delete more than two replication sets at a time, a job is created. You can follow the job’s progress and keep a history of its operation from the Job Summary pages.

Note: If you delete the last replication set in a consistency group, the consistency group is also deleted.

To discontinue replication and delete a replication set:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set summary page is displayed.

2. Select the replication set that you want to delete, and click Delete.

A confirmation message is displayed.

3. Click OK.

The replication set is deleted.

Note: You can also delete a replication set from its Details page by clicking the Delete button.

Adding a Replication Set to a Consistency Group

If you plan to add a replication set to a consistency group, the system administrator for the secondary peer must first make changes that match what you plan in order to ensure that the replication sets and consistency groups for the two peers continue to complement each other.

Note: The roles of the consistency group and the replication set must be the same. That is, a replication set must be the primary replication set if you plan to add it to a primary consistency group.

To add a replication set to an existing consistency group:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Click the consistency group to which you want to add the replication set.

The Consistency Group Details page is displayed.

3. Click Suspend and select Normal to place the consistency group into suspended mode.

4. From the navigation link under the navigational tabs, click Replication Set Summary.

The Replication Set Summary page is displayed.

5. Select a replication set that you want to add to the consistency group.

The Replication Set Details page is displayed.

6. Click Add To Group.
The Add To Consistency Group window is displayed.
7. Select the consistency group to which you want to add the replication set, and click OK.
8. Click the consistency group's name next to the Consistency Group label.
The Consistency Group Details page is displayed.
9. Click resume and select Normal synchronization to synchronize the volumes and resume replication.

Combining Replication Sets in a Consistency Group

If you have already created a number of replication sets and then determined that you want to place them in a consistency group, do so as outlined in the following sample procedure. In this example, Replication Set A and Replication Set B are existing independent replication sets. Follow these steps on both the primary and secondary peers:

1. Create a temporary volume, or identify an unused volume in the same storage domain as Replication Sets A and B.
2. Determine the World Wide Name (WWN) of the remote peer.
This information is on the Details page for either replication set.
3. Create Replication Set C. Do the following in the Create Replication Set wizard:
 - Select a temporary or unused volume from which to create the replication set.
 - In the Replication Peer WWN field, type the WWN of the remote system.
 - In the Remote Volume WWN field, type all zeros.
 - Create Consistency Group G.

Note: It is not necessary to create complementary replication sets on the remote peer at this time.
4. On the Details page for Replication Set A, click Add to Group to add the replication set to Consistency Group G.
5. On the Details page for Replication Set B, click Add to Group to add the replication set to Consistency Group G.

6. On the Details page for Replication Set C, click Delete to remove the replication set from Consistency Group G.

Replication Set A and Replication Set B are no longer independent and are now part of a consistency group.

Removing a Replication Set From a Consistency Group

If you plan to remove a replication set from a consistency group, the system administrator for the secondary peer must first make changes that match what you plan in order to ensure that the replication sets and consistency groups for the two peers continue to complement each other.

To remove a replication set from a consistency group:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Click the consistency group from which you want to remove the replication set.

The Consistency Group Details page is displayed.

3. Click Suspend, select Normal, and click OK to place the consistency group into suspended mode.

4. In the Replication Sets section, click the replication set that you want to remove from the consistency group.

The Replication Set Details page is displayed.

5. Click Remove From Group.

A confirmation message is displayed.

6. Click OK.

7. Click the consistency group's name next to the Consistency Group label.

The Consistency Group Details page is displayed.

8. Click Resume and select Normal synchronization to synchronize the volumes and resume replication.

Note: When you remove the last replication set from a consistency group, the consistency group is also deleted.

Deleting Consistency Groups

When you delete a consistency group, you also delete all the replication sets in the consistency group. The volumes in the consistency group and replication sets break their association and revert to conventional volumes.

To discontinue replication and delete a consistency group:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Click the consistency group that you want to delete.

The Consistency Group Details page is displayed.

3. Click Delete.

A confirmation message is displayed.

4. Click OK.

Reversing Roles

As a failover technique, you can reverse the roles of the primary and secondary peers in the event of a failure of the primary peer. The secondary peer can assume the role of the primary peer so that the application software now accesses the secondary volume directly while the failure of the primary peer is being addressed. You can reassign roles by changing the replication properties of a replication set or a consistency group.

The system administrator for the secondary peer must make changes that match what you plan in order to ensure that the replication properties of the replication sets and consistency groups for the two peers continue to complement each other.

To change a replication set's or consistency group's replication properties and reverse the roles of the peers:

1. Quiesce the application accessing the primary volume. Unmount the volume if necessary.

2. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

3. The Replication Set Summary page is displayed.

4. Select the replication set or consistency group whose role you want to change.

The Replication Set Details or Consistency Group Details page is displayed.

5. Click Suspend, select Normal, and click OK to place the replication set or consistency group into suspended mode.
6. In the Role Properties section of the Details page, select the new role from the Role menu for the replication set or consistency group.

7. Click Save.

The secondary peer now provides the original user data.

8. Click Resume and select Normal synchronization to synchronize the volumes and resume replication.
9. When the failure at the primary peer is corrected, you can reverse roles again so that both volumes return to their original roles.

Note: If a replication set a member of a consistency group, you cannot change the role of the replication set. The replication set's attributes must match those of the consistency group.

Note: If you use asynchronous mode during role reversal operations, the application might require extended down time while the asynchronous queue drains after you quiesce the application.

Changing Replication Modes

You must consider a number of factors and make a number of decisions before changing replication modes. For information on planning replication modes, see "About Replication Modes" on page 129 and "Reference: Synchronous Versus Asynchronous Replication Modes" on page 131.

You can change the replication mode of a replication set or a consistency group to synchronous or asynchronous mode.

To change the replication mode of a replication set or consistency group:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Click the replication set or consistency group whose replication mode you want to change.

The Replication Set Details or Consistency Group Details page is displayed.

3. Click Suspend and select Normal to place the replication set or consistency group into suspended mode.

4. Click the mode to which you want to change.
If you are changing to synchronous mode, a confirmation message is displayed. Click OK.
If you are changing to asynchronous mode, the Change Mode to Asynchronous wizard is displayed. Follow the steps in the wizard.
5. Click Resume and select Normal synchronization to synchronize the volumes and resume replication.

Replicating Data

To replicate data from the primary volume to a secondary volume:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.
The Replication Set Summary page is displayed.
2. Select the replication set or a consistency group that you want to replicate on the remote peer.
The Replication Set Details page or Consistency Group Details pages is displayed.
3. If you want synchronize the data on the volumes for both peers, follow these steps:
 - a. Click Resume.
The system displays the Resume Replication window.
 - b. Select the type of synchronization you want:
 - If you want to copy known differences rather than the entire volume, click Normal synchronization.
 - If you want to initiate a full volume-to-volume copy operation, click Full synchronization.
 - c. Click OK.
4. If you want to place the replication set or consistency group into suspended mode, click Suspend.
The system displays the Suspend Replication window.
5. If you want to suspend replication and track changes between the volumes in the replication bitmap, click Normal.
Note: If the replication set is in suspended mode, the Normal option is unavailable.

6. Click OK.
7. To resume replication, click Resume.

Synchronizing Data Using a Backup Tape

If you want to minimize data replication I/O traffic when you set up a copy of the data on a remote peer, you can use a backup tape copy of the primary volume to copy and synchronize data on the secondary volume.

To synchronize data using a backup tape:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

2. Select a replication set or a consistency group that you want to replicate to a remote peer.

The Replication Set Details page or Consistency Group Details pages is displayed.

3. Make sure that you have the autosynchronization option disabled so that you can start the synchronization operation manually when you have the backup tape ready.
4. Click Suspend and select Fast Start.
5. Have the system administrator for the secondary peer load the backup tape of the primary volume onto the secondary peer.
6. Click OK.
7. To resume replication, click Resume.

Troubleshooting Data Replication

If data replication does not occur, make sure that:

- The ports at both sites are enabled for data replication.
- You have an active replication link or connection.
- You have configured only two replication links, and the links are both on the same type of port. For example, both ports must be either Fibre Channel (FC) ports or Gigabit Ethernet ports, not a mix of the two.

- You entered the correct remote volume and peer port World Wide Names (WWNs) in the Create Replication wizard. Otherwise, you must delete the replication set and create a new one with the correct WWNs.
- Both peers have complementary replication set and consistency group definitions.
- The secondary volume includes snapshot reserve space if the secondary replication set has the Snap before Synchronization option enabled. For more information, see to “About Replication Set Properties” on page 126.

For information about FC and Gigabit Ethernet port failures, go to the Java Web Console and click Sun Storage Automated Diagnostic Environment.

Configuring Virtual Disks

This section describes virtual disks. It contains the following subsections:

- [“About Virtual Disks” on page 150](#)
- [“Displaying Virtual Disk Information” on page 151](#)
- [“Deleting a Virtual Disk” on page 151](#)

About Virtual Disks

A virtual disk is a set of disk blocks presented to an operating environment as a range of consecutively numbered logical blocks with disk-like storage and input/output semantics. The virtual disk is the disk array object that most closely resembles a physical disk from the operating environment’s viewpoint.

In the Sun StorEdge 6920 system, the system is considered the operating environment.

Redundant Array of Independent Disks (RAID) systems provide storage by making the data on many small disks readily available to file servers, hosts, or the network as a single array. RAID systems use two or more drives in combination for fault tolerance and performance. One of the factors in data throughput and availability is how the data is stored within the array, called the RAID level.

On the system, disk drives within a tray are grouped together into RAID sets, also called virtual disks, according to their RAID level. The storage arrays support the following RAID levels:

- **RAID-0** – Provides data striping across multiple disks, with no redundancy. This improves performance but does not deliver fault tolerance.
- **RAID-1** – Provides disk mirroring so that all data is copied to a separate disk. RAID-1 uses equal amounts of disk for the data and the mirror copy. RAID-1 can survive the loss of more than one disk, provided the disks aren’t copies of each other.
- **RAID-5** – Provides data striping at the block level as well as stripe error correction (parity checking) information. This is the system’s default RAID level. For this level, the minimum number of drives is three.

RAID-5 results in excellent performance and good fault tolerance. Parity checking specifies that when the RAID controller writes information onto disks, it also writes redundant information called parity bits. If a disk fails, the parity information enables the RAID controller to re-create the lost information as it is

requested. Since the parity information is spread across multiple disks, only a percentage of the disks is used for parity information. This improves the efficiency of available storage space.

Note: The maximum number of data partitions you can use per virtual disk is 32. A failure message is displayed if you try to exceed this number.

Use care in determining your storage pool and virtual disk utilization. A storage tray can have a maximum of two virtual disks. Therefore, if seven disks in a tray are assigned to a single storage pool, and you add one disk and configure it for a second storage pool, you cannot add six more disks later to the empty disk slots without reconfiguring the storage pools. In this situation, you cannot use the disks as virtual disks; however, you can use them as array hot-spares.

Displaying Virtual Disk Information

You can display summary and detailed information for existing virtual disks. You can also display summary information about disks and volumes associated with each virtual disk.

To display information for virtual disks:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Virtual Disks.
The Virtual Disk Summary page is displayed.
2. Select a virtual disk to see its details.
The Virtual Disk Details page is displayed.
3. Go to Additional Information to see the number of disks, replication sets, and volumes associated with the selected virtual disk.
Click an item to see its Summary page.

Deleting a Virtual Disk

You cannot delete a virtual disk while it is in use.

To delete a virtual disk:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Virtual Disks.
The Virtual Disk Summary page is displayed.

2. Select the check box to the left of the virtual disk you want to delete.
3. Click Delete.

A confirmation message displays.

4. Click OK.

A confirmation message is displayed, and the virtual disk is removed from the Virtual Disk Summary page.

Note: To delete multiple virtual disks at once, select the check box for each virtual disk you want to delete, and then click Delete.

Configuring Storage Pools

This section describes storage pools. It contains the following subsections:

- [“About Storage Pools” on page 153](#)
- [“Planning Storage Pools” on page 153](#)
- [“Managing Storage Pools” on page 154](#)

About Storage Pools

The default configuration uses the DEFAULT storage pool, which consists of all the virtual disks in the system. When you create a new storage pool, you assign virtual disks to it, and then you create volumes of disks.

Each storage pool is assigned a storage profile that meets specific I/O requirements. The system provides a set of storage profiles that satisfy certain I/O requirements, but if none of the profiles are suitable for your site’s need, you can create a custom storage profile. For more information, see [“Creating a Storage Profile” on page 165](#).

To see the current storage pools, go to the Storage Pool Summary page, as described in [“Displaying Pool Information” on page 154](#).

Planning Storage Pools

The number of storage pools you need for a storage domain depends on your array configuration. Each storage tray can support up to two virtual disks (with one virtual disk in each storage pool); for example, in a 2x6 array configuration, you can define a maximum of 12 storage pools.

Use care in determining your storage pool and virtual disk utilization. A storage tray can have a maximum of two virtual disks. Therefore, if seven disks in a tray are assigned to a single storage pool, and you add one disk and configure it for a second storage pool, you cannot add six more disks later to the empty disk slots without reconfiguring the storage pools. In this situation, you cannot use the disks as virtual disks; however, you can use them as array hot-spares. For more information about virtual disks, see [“About Virtual Disks” on page 150](#).

Before adding a new storage pool, determine the following:

- The storage domain with which you will associate the pool

- The storage profile with which you will associate the pool
The attributes of the storage profile will determine the storage characteristics of the virtual disks you can configure for the new storage pool.
Note: You create virtual disks from storage trays (by adding storage to the pool) or from external (raw) storage.
- The number of disks and the target size for the pool, if you want to add storage to the pool when you create it.
The system will provide a list of available capacities, based on the configuration method and target size you choose, from which you specify the number of blocks, Kbytes, Mbytes, Gbytes, or Tbytes.

Reconfiguring a storage pool consists of removing the storage pool and then creating a new storage pool. Storage pools can be expanded dynamically so you can remove unused virtual disks and add new virtual disks.

Note: You cannot delete a storage pool to which data is allocated, until you first delete the volumes and move, remove, or delete the virtual disks within the pool.

Managing Storage Pools

This section describes ways in which you can manage storage pools. It contains the following subsections:

- [“Displaying Pool Information” on page 154](#)
- [“Creating a Storage Pool” on page 155](#)
- [“Adding Storage to a Pool” on page 155](#)
- [“Modifying a Storage Pool” on page 156](#)
- [“Deleting a Storage Pool” on page 156](#)

Displaying Pool Information

You can display summary information and details about storage pools that exist on your system, and you can view additional information about the virtual disks and volumes associated with each storage pool.

To display information on pools:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Pools.

The Storage Pool Summary page is displayed.

2. Select a storage pool to see its details.

The Storage Pool Details page is displayed.

3. Go to Additional Information to see virtual disks and volumes associated with the selected storage pool.

Click an item to display its Summary page.

Creating a Storage Pool

You must consider a number of factors and make a number of decisions before creating a storage pool. For information on planning a storage pool, see “Planning Storage Pools” on page 153.

To create a storage pool:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Pools.

The Storage Pool Summary page is displayed.

2. Click New.

The New Storage Pool wizard is displayed.

3. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Adding Storage to a Pool

You must consider a number of factors and make a number of decisions before adding storage to a pool. For information on planning a storage pool, see “Planning Storage Pools” on page 153.

When configuring a storage pool, you add storage to the pool so you can later store data.

To add storage to a pool:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Pools.

The Storage Pool Summary page is displayed.

2. Select the storage pool to which you want to add storage.

The Storage Pool Details page is displayed.

3. Click Add Storage to Pool.

The Add Storage wizard is displayed.

4. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Modifying a Storage Pool

To modify a storage pool:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Pools.

The Storage Pool Summary page is displayed.

2. Select the storage pool that you want to modify.

The Storage Pool Details page is displayed.

3. Modify the name and/or description for the selected storage pool, and click Save.

Deleting a Storage Pool

You cannot delete storage pools that have volumes/virtual disks associated with them. For instructions on deleting a volume, see “Deleting a Volume” on page 90. For information on deleting a virtual disk, see “Deleting a Virtual Disk” on page 151.

To delete a storage pool:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Pools.

2. Select the check box to the left of each storage pool that you want to delete.

This enables the Delete button.

3. Click Delete.

A confirmation message is displayed

4. Click OK.

The selected storage pool no longer appears in the Storage Pools table.

Configuring Storage Profiles

This section describes storage profiles. It contains the following subsections:

- [“About Storage Profiles” on page 158](#)
- [“About Virtualization Strategies” on page 161](#)
- [“Reference: Standard Storage Profiles” on page 162](#)
- [“Managing Storage Profiles” on page 164](#)

About Storage Profiles

A storage profile is set of attributes that you apply to a storage domain to allocate storage, instead of having to set each attribute individually. The system has a predefined set of storage profiles, including the Default storage profile. You can choose a storage profile suitable for the application that is using the storage, or you can create a custom storage profile.

The following table lists and describes the parameters for storage profile settings.

TABLE 4-5 Parameters for Storage Profile Settings

Parameter	Value or Variable Type	Description
Name	Up to 32 characters	Unique identifier for the storage profile.
Description	Up to 256 characters	Description of the storage profile.
RAID Level	The RAID level of the virtual disk: <ul style="list-style-type: none">• RAID-0• RAID-1• RAID-5	RAID level configured across all disks within a virtual disk.

TABLE 4-5 Parameters for Storage Profile Settings (*Continued*)

Parameter	Value or Variable Type	Description
Segment Size	4 KB, 8 KB, 16 KB, 32 KB, or 64 KB	<p>The size of the array segment in kilobytes: 4 KB, 8 KB, 16 KB, 32 KB, or 64 KB.</p> <p>This is an array attribute and must be the same for all virtual disks within an array. Once a virtual disk exists on any of the array's trays (there can be only 0, 1, or 2 virtual disks on a tray), then all of the virtual disks on that array's trays must have the same segment size as that virtual disk. Be careful when allocating segment size to a profile, as you cannot have pools whose profiles differ in segment size.</p> <p>You can configure a system with an array that has unused disks in which no virtual disks can be configured if the pools' profiles do not have the same segment size.</p>
Dedicated Hot-Spare	Enabled or Disable	<p>Whether each virtual disk requires an additional disk from the same tray to be used as its dedicated hot-spare.</p> <p>The system has dedicated hot-spares (controlled by this attribute) and array hot-spares (controlled as an editable attribute on the Array Details page).</p> <ul style="list-style-type: none">• Dedicated hot-spare – There is a one-to-one correspondence between a dedicated hot-spare and a virtual disk (that is, the hot-spare is dedicated to a specific virtual disk). This profile attribute controls whether a dedicated hot-spare disk is assigned to the virtual disk when a virtual disk is created in a pool with this profile.• Array hot-spare – This is a disk in an array that can act as a hot-spare for any disk in the array. The number of drives and dedicated hot-spare attributes taken together define how many disks are consumed for each virtual disk in a pool.

TABLE 4-5 Parameters for Storage Profile Settings (*Continued*)

Parameter	Value or Variable Type	Description
Read Ahead	On (Enabled) or Off (Disabled)	<p>Read-ahead mode of the array. This array attribute must be the same for all virtual disks within an array. Once a virtual disk exists on any of an array's trays (you can have only 0, 1, or 2 virtual disks on a tray), all of the virtual disks on that array's trays must have a Read-Ahead Enabled setting that matches that virtual disk's Read-Ahead Enabled setting.</p> <p>You must plan the use of the Read-Ahead Enabled setting carefully, because you cannot have pools whose profiles have different Read-Ahead Enabled settings on the same array. You can configure a system with unused disks in which no virtual disks can be configured if the pools' profiles do not match the Read-Ahead Enabled setting.</p>
# Drives	Variable, 2 through 14	<p>The number of drives associated with this profile. This is the number of disk drives that are used to construct the virtual disk in a pool using this profile.</p> <p>The number of drives can be 2 to 14, or Variable. If it is set to variable, you must specify the number when you create virtual disks (that is, when you add storage to the pool).</p>
Array Type	Exact Match: IOPS (2x4 array) Exact Match: Bandwidth (2x2 array) Exact Match: Capacity (2x6 array) Best Available Match: IOPS Best Available Match: Bandwidth Best Available Match: Capacity	<p>The Exact Match options must be used only with the specified array configuration (2x2, 2x4, and 2x6) so, for a single system, only one of the Exact Match options is appropriate. If a profile defines an Exact Match option but is applied to the wrong array type, you cannot add storage to that pool.</p> <p>The Best Available options attempt to match the array type based on IOPS (input/output performance speed of transactions per second), bandwidth, or capacity.</p> <p>The following rules apply: Bandwidth: 2x2 > 2x4 > 2x6 Capacity: 2x6 > 2x4 > 2x2 IOPS: 2x4 > 2x2 > 2x6</p>

TABLE 4-5 Parameters for Storage Profile Settings (*Continued*)

Parameter	Value or Variable Type	Description
Virtualization Strategy	Concatenate or Stripe	Strategy to be used for all volumes allocated from a storage pool. The Concatenate value: <ul style="list-style-type: none">• Requires a minimum of one virtual disk• Is used in sequential environments• Implements algorithms to create volumes from “best fit” virtual disks to minimize overall partitioning of virtual disks The Stripe value: <ul style="list-style-type: none">• Requires a minimum of two virtual disks• Is used in random environments• Stripes across multiple partitions in volumes
Stripe Size	128 KB, 512 KB, 1 MB, 2 MB, 4 MB, 8 MB, or 16 MB	The amount of space in memory on a virtual disk to be allocated for the stripe used in the Stripe virtualization strategy. This field is valid when the Virtualization parameter is set to Stripe.

About Virtualization Strategies

Before you create a storage profile, determine the virtualization strategy you want to associate with the storage profile.

You can choose between two virtualization strategies:

- **Concatenation** – In a concatenated volume, virtual disk partitions are filled with data sequentially, with the second disk being written to when no space remains on the first, the third when no room remains on the second, and so on. The system distributes the volume space across devices by using the devices with the smallest amount of available space first.
- **Striping** – In a striped volume, the system divides data into blocks and distributes (or stripes) the data blocks evenly across several disk blocks. The system distributes the volume space across devices equally by using the smallest number of devices possible to make a volume of the desired size. Striping distributes (reads and writes) data to two or more slices that reside on different disks. Striping data improves your disk performance because it performs simultaneous disk read/write operations on multiple disks.

Reference: Standard Storage Profiles

The following table describes the storage profiles included in the Sun StorEdge 6920 system.

TABLE 4-6 Standard Storage Profiles in the Sun StorEdge 6920 System

Profile Name	RAID Level	Segment Size	Read Ahead Mode	Number of Drives	Array Type	Dedicated Hot-Spare	Virtualization	Stripe Size
Default	RAID-5	64 KB	On	Variable	Best Available Match: Bandwidth (2x2 > 2x4 > 2x6)	No	Stripe	1 MB
Sequential	RAID-5	64 KB	On	Variable	Exact Match: Bandwidth (2x2 array)	No	Concatenate	n/a
Random 5	RAID-5	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB
Random 1	RAID-1	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB
High Performance Computing	RAID-5	64 KB	On	Variable	Exact Match: Bandwidth (2x2 array)	No	Concatenate	n/a
Mail Spool	RAID-1	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB

TABLE 4-6 Standard Storage Profiles in the Sun StorEdge 6920 System (Continued)

Profile Name	RAID Level	Segment Size	Read Ahead Mode	Number of Drives	Array Type	Dedicated Hot-Spare	Virtualization	Stripe Size
NFS Striping	RAID-5	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB
NFS Mirroring	RAID-1	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB
Oracle OLTP	RAID-5	32 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	1 MB
Oracle OLTP HA	RAID-1	32 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	1 MB
Oracle DSS	RAID-5	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB
Sybase OLTP	RAID-5	32 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	1 MB

TABLE 4-6 Standard Storage Profiles in the Sun StorEdge 6920 System (Continued)

Profile Name	RAID Level	Segment Size	Read Ahead Mode	Number of Drives	Array Type	Dedicated Hot-Spare	Virtualization	Stripe Size
Sybase OLTP HA	RAID-1	32 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	1 MB
Sybase DSS	RAID-5	64 KB	On	Variable	Best Available Match: IOPS (2x4 > 2x2 > 2x6)	No	Stripe	2 MB

Managing Storage Profiles

This section describes ways in which you can manage storage profiles. It contains the following subsections:

- [“Displaying Profile Information” on page 164](#)
- [“Creating a Storage Profile” on page 165](#)
- [“Modifying a Storage Profile” on page 165](#)
- [“Deleting a Storage Profile” on page 166](#)

Displaying Profile Information

You can display summary and detail information about existing storage profiles. You can also display summary information about pools and volumes associated with each storage profile.

To display information about storage profiles:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
2. Select a profile to see its details.
The Storage Profile Details page is displayed.

3. Go to Additional Information and click Storage Pools or Volumes for additional information associated with the selected profile.

The Summary page for the selected item is displayed.

Creating a Storage Profile

You must make a number of decisions about profile characteristics before you create a profile. For information on planning a profile, see “About Storage Profiles” on page 158.

To create a storage profile:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Profiles.

The Storage Profile Summary page is displayed.

2. Click New.

The New Storage Profile wizard is displayed.

3. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Modifying a Storage Profile

You cannot modify the default factory profiles. If a user-created profile’s state is In Use, you can change only the profile name and description. A profile is in the In Use state when it is associated with a storage pool.

To modify a storage profile:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Profiles.

The Storage Profile Summary page is displayed.

2. Select the storage profile that you want to modify.

The Storage Profile Details page is displayed.

3. Make the appropriate modifications, and click Save.

Deleting a Storage Profile

You cannot delete the Default storage profile, nor can you delete a profile that is in the In Use state and associated with a storage pool.

To delete a storage profile:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Profiles.

The Storage Profile Summary page is displayed.

2. Select the check box for the storage profile that you want to delete. This enables the Delete button.

Note: You can multiple check boxes to delete multiple storage profiles.

3. Click Delete.

The selected storage profile is removed from the Storage Profile Summary page.

Configuring Storage Domains

This section describes storage domains. It contains the following subsections:

- [“About Storage Domains” on page 167](#)
- [“Managing Storage Domains” on page 169](#)

About Storage Domains

One unique design element of the Sun StorEdge 6920 system is the concept of storage domains. Storage domains are “containers” of physical devices and virtual (logical) storage elements managed by the system.

A storage domain includes a virtual (internal) network, connections to physical devices, and policies that govern the storage domain.

- The physical devices include internal storage arrays, external arrays, and servers.
- The virtual storage elements include virtual disks, storage profiles, storage pools, and virtual volumes.

Storage domains allow data to be isolated so that data from one host or set of hosts is protected from access by other hosts. In its default configuration, the system is provisioned with one storage domain, called DEFAULT. This storage domain contains the Default storage pool with its associated disks and virtual disks.

To group servers together and provision storage for that group of servers only, you redefine the DEFAULT storage domain into two or more storage domains. A configuration with multiple storage domains provides an environment in which multiple servers share storage without compromising client security, as shown in the following figure.




FIGURE 4-4 Configuration With Multiple Storage Domains

You can create a maximum of 12 storage domains within one system. Within a storage domain, you can have multiple storage pools. However, each storage pool can be associated with only one storage domain. For example, suppose the Finance domain contains the Random I/O and Mail Space pools, the Engineering domain contains the File Serving pool, and the Sales domain contains the Sequential pool. In this case, the Finance domain cannot also be associated with the File Serving or Sequential pool, the Engineering domain cannot also be associated with the Random I/O, Mail Space, or Sequential pool, and the Sales domain cannot be also associated with the Random I/O, Mail Space, or File Serving pool.

To see the current storage domains, go to the Storage Domain Summary page, as described in “Displaying Storage Domain Information” on page 169.

Managing Storage Domains

This section describes ways in which you can manage storage domains. It contains the following subsections:

- [“Displaying Storage Domain Information” on page 169](#)
- [“Creating a Storage Domain” on page 169](#)
- [“Deleting a Storage Domain” on page 170](#)

Displaying Storage Domain Information

You can display summary and detailed information about your system’s storage domains, and you can view information about initiators, storage pools, and volumes associated with each storage domain.

To display information about storage domains:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Domains.
The Storage Domain Summary page is displayed.
2. Select a storage domain name for detailed information.
The Storage Domain Details page for the selected domain is displayed.
3. Go to Additional Information to see the number of initiators, storage pools, and volumes associated with the selected storage domain.
Click an item to see its summary page.

Creating a Storage Domain

Before you begin, decide which initiators you want to assign to the new storage domain.

To create a storage domain:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Domains.
The Storage Domain Summary page is displayed.
2. Click New.
The Create Storage Domain wizard is displayed.

3. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Deleting a Storage Domain

You can delete a domain only if there are no volumes and no pools associated with it.

Note: You cannot delete the DEFAULT storage domain.

To delete a storage domain:

1. Click Sun StorEdge 6920 Configuration Service > Logical Storage > Domains.

The Storage Domain Summary page is displayed.

2. Select the domain that you want to delete.

The Storage Domain Details page is displayed.

3. Click Delete to remove the storage domain from your system.

Configuring Initiators

This section describes initiators. It contains the following subsections:

- [“About Initiators” on page 171](#)
- [“Managing Initiators” on page 171](#)

About Initiators

An initiator is the protocol-specific physical port that initiates the I/O exchanges with the system. In a Fibre Channel (FC) storage area network (SAN), an initiator is the FC port that is identified by a port World Wide Name (WWN). If a host is connected to the system by two host bus adapters (HBAs), the system handles the host as two different initiators. FC array Logical unit number (LUN) masking and mapping uses initiator port identifiers to authenticate storage customers.

In a Small Computer System Interface (SCSI) environment, an initiator is a SCSI node that sends requests to a SCSI device. SCSI adapters in hosts are usually initiators, but storage devices that provide virtualization can also act as initiators to other SCSI devices.

When a new initiator is connected into the storage environment, the system discovers it and adds it to the DEFAULT storage domain. To complete the configuration, add a description of the initiator to identify it, move it to another storage domain if necessary, and then map the initiator to a volume.

When you create a new storage domain, you specify which unmapped initiators are in the domain. After an initiator is mapped to a volume, you cannot change its domain.

To see the current initiators, go to the Initiator Summary page, as described in [“Displaying Initiator Information” on page 172](#).

Managing Initiators

This section describes ways in which you can manage initiators. It contains the following subsections:

- [“Displaying Initiator Information” on page 172](#)
- [“Creating an Initiator” on page 172](#)

- [“Mapping an Initiator to Volumes” on page 173](#)
- [“Deleting an Initiator” on page 173](#)
- [“Unmapping an Initiator From a Volume” on page 174](#)

Displaying Initiator Information

You can display summary and detail information about existing initiators. You can also display summary information about mapped volumes associated with an initiator you select.

To display information about initiators:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Initiators.
The Initiator Summary page is displayed.
2. Select an initiator for detailed information.
The Initiator Details page for the selected initiator is displayed.
3. Go to Additional Information to see the number of mapped volumes associated with the selected initiator.
Click Mapped Volumes to see its Summary page.

Creating an Initiator

You can create a new initiator and then associate it with a volume. Alternatively, you can configure an initiator before a physical connection has been established. In this case, you might set up the volumes, map them to the initiator, and then provide the World Wide Name (WWN).

You will need the WWN for the initiator you want to associate with a volume.

To create an initiator:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Initiators.
The Initiator Summary page is displayed.
2. Click New.
The New Initiator wizard is displayed.
3. Follow the steps in the wizard.
4. Click the Help tab in the wizard for more information.

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Mapping an Initiator to Volumes

To map initiators to a volume:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Initiators.
The Initiator Summary page is displayed.
2. Select the check box for the initiator you want to map to a volume.
This enables the Map button.
Note: You can select multiple initiators to map to a volume.
3. Click Map.
4. The Map Initiators wizard is displayed.
5. Follow the steps in the wizard.
Click the Help tab in the wizard for more information

Note: If you navigate to another tab within the browser interface while you are following steps in a wizard, a blank page may appear when you finish and close the wizard. If this occurs, you can close the blank page without danger of losing any data.

Deleting an Initiator

You cannot delete an initiator that is mapped to a volume (In Use state).

To delete an initiator:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Initiators.
The Initiator Summary page is displayed.
2. Select the check box for the initiator that you want to delete.
This enables the Delete button.
Note: You can select multiple initiators to delete.

3. Click Delete.

The selected initiators are removed from the Initiator Summary page.

Unmapping an Initiator From a Volume

To unmap initiators from a volume:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Initiators.

The Initiator Summary page is displayed.

2. Select the initiator you want to unmap.

The Initiator Details page is displayed.

3. Go to the Additional Information table and select Mapped Volumes.

The Volume Mapping Summary page is displayed.

4. Click the volumes you want to unmap from the initiators, then click Unmap.

5. Click OK to confirm.

The initiators are unmapped from the selected volumes.

Configuring System Ports

This section describes viewing and enabling Fibre Channel (FC) and Gigabit Ethernet ports. It contains the following subsections:

- [“About Ports” on page 175](#)
 - [“Displaying Port Information” on page 176](#)
 - [“Displaying Replication Link Information” on page 177](#)
 - [“Enabling Ports for Data Replication” on page 177](#)
 - [“Disabling Ports for Data Replication” on page 178](#)
 - [“Adding Replication Links” on page 179](#)
 - [“Deleting Replication Links” on page 179](#)
-

About Ports

The system’s service processor panel includes Fibre Channel (FC) and Gigabit Ethernet ports. The number of available ports is determined by the number of storage resource card (SRC) sets installed in the system’s Data Services Platform (DSP).

FC ports are connected to data hosts, external storage devices, or another Sun StorEdge 6920 system. These are the outbound ports from the service panel to the data hosts, external storage devices, or other Sun StorEdge 6920 system, and the inbound ports from the DSP to the system arrays.

You can use the system’s data replication feature to replicate data between your system and another Sun StorEdge 6920 system. You can use either FC or Gigabit Ethernet ports for data replication; you cannot use both a Gigabit Ethernet port and a FC port on a system to replicate data. Gigabit Ethernet ports allow you to connect to another Sun StorEdge 6920 system over a wide area network (WAN) for purposes of data replication.

Note: Your system must have a Gigabit Ethernet port connection if you plan to use the data replication feature over IP.

Before you can replicate data from your system to a remote Sun StorEdge 6920 system, you must:

- Make sure that the remote system is available to your system over an existing link.

- If you are using FC ports, configure the FC switch you use to make the connection between sites for long-distance operations. See the FC switch vendor's documentation for information about operating over long distances.
- Enable the ports on both systems for replication.

Note: You must enable the ports before you can proceed to the task to prepare for data replication.

If you enable a Gigabit Ethernet port for data replication, complete the Enable Gigabit Ethernet Port Replication wizard to configure the remote system's IP information on the Gigabit Ethernet port and create a replication link.

A replication link is the logical interconnection between two Sun StorEdge 6920 systems that carries control information and data from the local system (or primary host) to the remote system (or secondary host).

Before you use the wizard to configure the Gigabit Ethernet ports for replication, you must have the following information:

- The default gateway, network mask, and address of the Gigabit Ethernet port on your system.
- The IP address of the Gigabit Ethernet ports on the remote system.

The wizard allows you to specify IP addresses of up to four remote peer systems. You can specify additional replication links from the Gigabit Ethernet Port Details page. For more information, see "Adding Replication Links" on page 179.

The system administrator for the remote system must also take the above considerations into account, enabling ports, and creating replication links to ensure connectivity between the two systems.

Refer to "Managing Data Replication" on page 139 for more information about the data replication feature.

Displaying Port Information

You can display summary and detail information about existing Fibre Channel (FC) and Gigabit Ethernet ports. You can also display summary information about the initiators that are associated with each FC port and the replication links that are associated with each Gigabit Ethernet port.

To display information about ports:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.

The Port Summary page is displayed.

2. Click a Port name for detailed information about that port.
The Port Details page for the selected port is displayed.
3. Go to Additional Information and click any item for more information associated with the selected port.
The Summary page for the selected item is displayed.

Displaying Replication Link Information

You can display information about replication links that are associated with a Gigabit Ethernet port for use with the data replication feature.

To display information about replication links:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.
The Port Summary page is displayed.
2. Click a Gigabit Ethernet port name.
The Port Details page for the selected Gigabit Ethernet port is displayed.
3. Go to Additional Information and click Replication Links.
The Replication Link Summary page is displayed.

Enabling Ports for Data Replication

You must consider a number of factors and make a number of decisions before using a Gigabit Ethernet port for data replication. For information on planning data replication, see “Planning for Data Replication” on page 134.

To enable a Fibre Channel (FC) or Gigabit Ethernet port for data replication:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.
The Port Summary page is displayed.
2. Click a port name.
The Fibre Channel Port Details or Gigabit Ethernet Port Details page is displayed, depending on the type of port you select.

3. If the port is an FC port, click OK to confirm that you want to enable the port for replication.

The port's replication status displays as Enabled.

4. If the port is a Gigabit Ethernet port, click Configure Replication.

The Configure Gigabit Ethernet Port Replication wizard is displayed.

5. Complete the steps in the wizard.

Disabling Ports for Data Replication

When you disable ports for data replication, the impact depends on how you configured your system. If you have two ports enabled for data replication and you disable one, the available data replication bandwidth is halved. If you have multiple ports enabled for data replication and you disable one, you might not notice decrease in data replication.

To disable a Fibre Channel (FC) or Gigabit Ethernet port for data replication:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.

The Port Summary page is displayed.

2. Click a Port name.

The Fibre Channel Port Details or Gigabit Ethernet Port Details page is displayed, depending on the type of port you select.

3. Click Disable Replication.

The system requires you to provide confirmation.

4. Click OK to confirm that you want to disable the port for replication.

The port's Replication status displays as Disabled.

Note: When you disable a Gigabit Ethernet port, all the replication information associated with that port and the replication links is deleted.

Adding Replication Links

To add a replication link:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.
The Port Summary page is displayed.
2. Click a Gigabit Ethernet port name.
The Port Details page for the selected port is displayed.
3. Go to Additional Information and click Replication Links.
The Replication Link Summary page is displayed.
4. Click Add Link.
The Add Replication Link page is displayed.
5. Type the IP address of a remote peer system and click OK.

Deleting Replication Links

To delete a replication link:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Ports.
The Port Summary page is displayed.
2. Click a Gigabit Ethernet port name.
The Port Details page for the selected port is displayed.
3. Go to Additional Information and click Replication Links.
The Replication Link Summary page is displayed.
4. Click Remove Link.
The system queries if you want to delete the link.
5. Confirm that you want to delete the replication link.

Configuring Storage Arrays

This section describes storage arrays. It contains the following subsections:

- [“About Arrays and Trays” on page 180](#)
 - [“About Conventions for Array Names and IP Addresses” on page 180](#)
 - [“Displaying Storage Array Information” on page 182](#)
 - [“Modifying an Array” on page 183](#)
-

About Arrays and Trays

Because trays are always configured in pairs for high availability (HA), each HA array configuration contains 2 controller trays and from 0 to 4 expansion trays, for a total of up to 6 trays and a maximum of 84 disks per array.

- The number of storage trays that can be installed in a cabinet is limited by the space available in the cabinet and by the array configuration.
- The number of arrays is limited by the number of storage ports available. Each array is connected to a single storage port through a Fibre Channel (FC) port connection on the service panel.
- For an expansion cabinet, the number of arrays is also limited by the number of arrays installed in the base cabinet. In a second expansion cabinet, the number of arrays is limited by the number of arrays in both the first and the second cabinets.

To see the current storage trays, go to the Storage Tray Summary page, as described in [“Displaying Storage Tray Information” on page 188](#).

To see the current storage arrays, go to the Storage Array Summary page, as described in [“Displaying Storage Array Information” on page 182](#).

About Conventions for Array Names and IP Addresses

The host names and Internet Protocol (IP) addresses for arrays are based on the arrays' positions in the cabinet:

- Each array has a static, preconfigured IP address. The array in the lowest position in a base cabinet has the IP address 192.168.0.40.

- Each array has a host name formed by the following syntax:

array + cabinet number + position number

where:

- *cabinet number* is 0 for the base cabinet, 1 for the first expansion cabinet, 2 for the second expansion cabinet, and so on.
- *position number* is 0 through 3 for positions in the base cabinet and 0 through 5 for positions in an expansion cabinet, 0 being the lowest position.

The array in the lowest position in the base cabinet has the host name `array00`, indicating that it is the array on the bottom of cabinet 0. The lowest array in the first expansion cabinet is `array10`, and the lowest array in the second expansion cabinet is `array20`.

Within each array, the bottom controller tray is the master and the top controller tray is the alternate master. For example, the controller tray in `tray1` is the master, and the controller tray in `tray2` is the alternate master. If the master controller tray fails, the alternate master controller tray uses the array's IP address.

The following figure shows an overview of how storage trays and arrays are designed for a typical Sun StorEdge 6920 system.




FIGURE 4-5 How Storage Trays and Arrays Are Designed for a Typical Sun StorEdge 6920 System

Displaying Storage Array Information

You can display summary and detail information for a storage array. You can also display summary information about the disks, trays, virtual disks, and volumes associated with each array.

To display information for a storage array:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Arrays.
The Array Summary page is displayed.

2. Select the array you want to view.

The Array Details page is displayed.

3. Go to Additional Information to see the number of disks, replication sets, trays, virtual disks, and volumes associated with the selected array.

Click an item to display its Summary page.

Modifying an Array

You can modify the number of array hot-spares, disk reconstruction rate, cache mode, and diagnostic mode, and you can enable or disable the disk scrubbing option for an array.

To modify an array:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Arrays.

The Array Summary page is displayed.

2. Select the array you want to modify.

The Array Details page for the selected array is displayed.

3. Modify any of the following array fields:

- Array Hot Spares: 0 or n
- Disk Reconstruction Rate: Medium (default), Low, High
- Cache Mode: Auto (default), WriteBehind, WriteThrough, Disabled
- Diagnostic Mode: BEFIT, Off, Active, Passive
- Disk Scrubbing: Enabled or Disabled

4. Click Save.

Using Storage Trays

This section describes storage trays. It contains the following subsections:

- [“About Storage Trays” on page 184](#)
- [“About Tray Numbering Conventions” on page 185](#)
- [“About Identifiers for Cabinets” on page 186](#)
- [“Displaying Storage Tray Information” on page 188](#)

About Storage Trays

Storage trays contain physical disk drives. Each tray has 7 to 14 disk drives and is identified by whether it contains a Redundant Array of Independent Disks (RAID) controller:

- A **controller tray** is a storage tray with hardware RAID management built into the controller.
- An **expansion tray** is a storage tray with disk drives only and no controller. This type of tray is managed by a controller tray.

The following figure shows the symbols for a controller tray and an expansion tray.




FIGURE 4-6 Symbols for a Controller and an Expansion Tray

The following table shows how the trays are arranged in the supported array configurations.

TABLE 4-7 How Trays are Arranged in the Supported Array Configurations

Array Configuration	Number of Controller Trays	Number of Expansion Trays
2x2	2	0
2x4	2	2
2x6	2	4

The following figure shows how the trays can be configured.




FIGURE 4-7 How Trays Can Be Configured

To see the current storage trays, go to the Storage Tray Summary page, as described in “Displaying Storage Tray Information” on page 188.

To see the current storage arrays, go to the Storage Array Summary page, as described in “Displaying Storage Array Information” on page 182.

Use care in determining your storage pool and virtual disk utilization. A storage tray can have a maximum of two virtual disks. Therefore, if seven disks in a tray are assigned to a single storage pool, and you add one disk and configure it for a second storage pool, you cannot add six more disks later to the empty disk slots without reconfiguring the storage pools. In this situation, you cannot use the disks as virtual disks; however, you can use them as array hot-spares. For more information about virtual disks, see “About Virtual Disks” on page 150.

About Tray Numbering Conventions

Trays are numbered from the bottom to the top, starting with 1.

The following figure shows the tray numbers and where the controller trays are located in each possible array configuration.

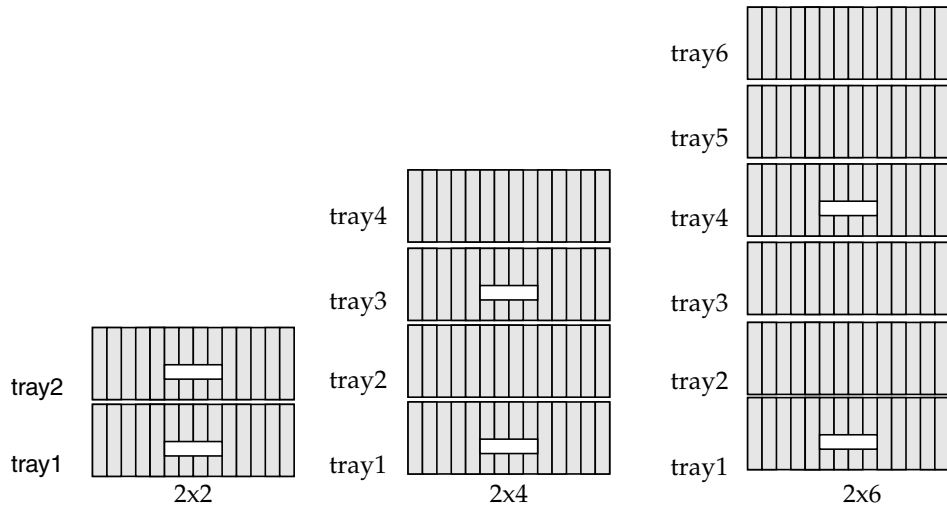


FIGURE 4-8 Tray Numbers and Controller Tray Locations

Each array has a designated master controller tray and an alternate master controller tray. If the master controller tray fails, the alternate master assumes control of the array's operations. The master controller tray is the lowest controller tray of the two until the first failover, as shown in the following table.

TABLE 4-8 Tray Master and Alternate Master

Tray	2x2	2x4	2x6
Master	tray1	tray 1	tray1
Alternate master	tray2	tray3	tray4

For information on array host names and IP addresses, see "About Conventions for Array Names and IP Addresses" on page 180.

About Identifiers for Cabinets

Each system consists of a base cabinet and, optionally, up to two expansion cabinets. The cabinet numbers are used in the names and Internet Protocol (IP) addresses of the arrays and of other hardware in the cabinet.

The cabinets are numbered from 0 through 2, as follows:

- Base cabinet: 0

- Expansion cabinet: 1
- Expansion cabinet: 2

The maximum number of trays in the base cabinet is 8. The maximum number of trays in an expansion cabinet is 12. The following figure shows the maximum number of trays and how the trays are numbered. (Each tray is shown as a controller tray.)

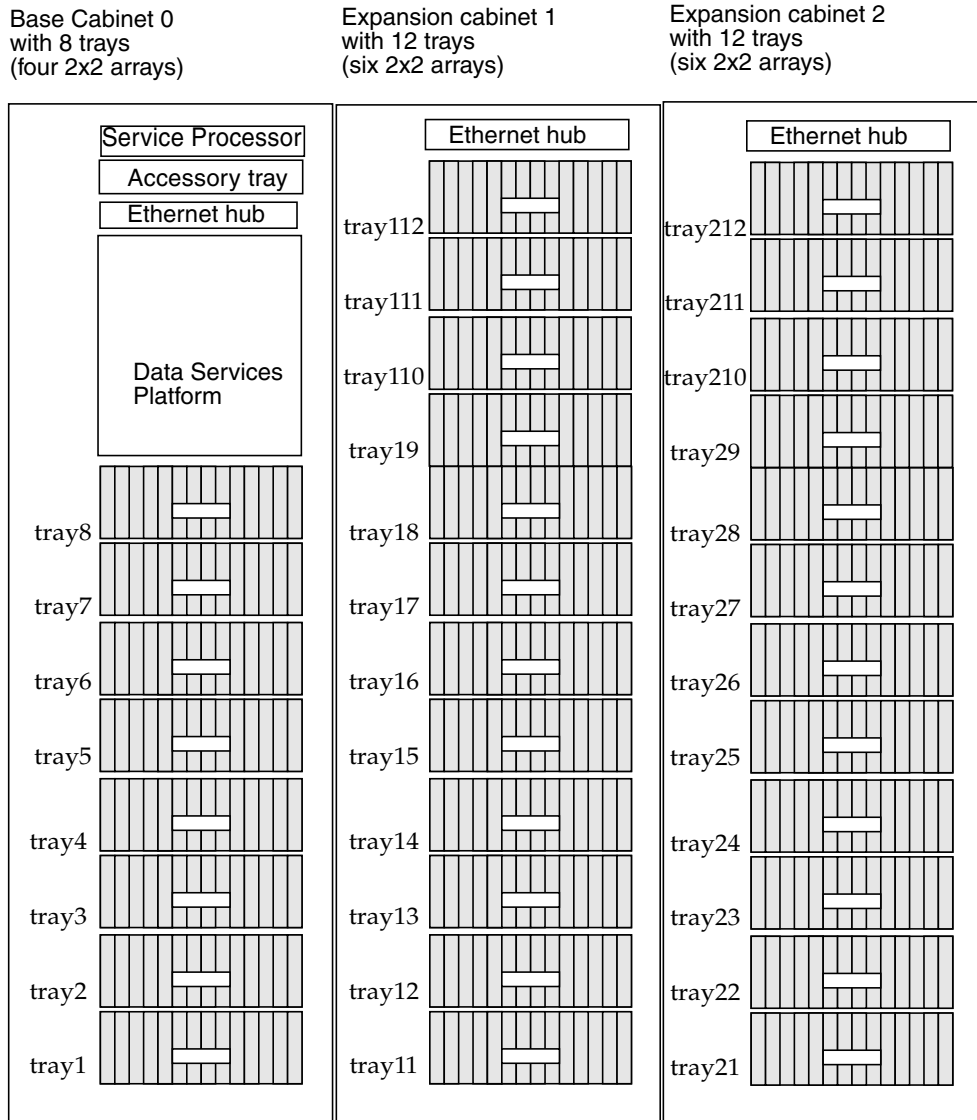


FIGURE 4-9 Maximum Number of Trays and How They Are Numbered

Displaying Storage Tray Information

You can display summary and detail information about existing storage trays. You can also display summary information about the disk drives, virtual disks, and volumes that are associated with each tray.

To display information on trays:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Trays.

The Tray Summary page is displayed.

2. Select the tray you want to view.

The Tray Details page is displayed.

3. Click Additional Information to view the number of disks, replication sets, virtual disks, and volumes associated with the selected tray.

Click an item to see its Summary page.

Managing Disk Drives

This section describes disk drives. It contains the following subsections:

- [“About Disk Drives” on page 189](#)
- [“Displaying Disk Information” on page 190](#)

About Disk Drives

The following table describes the disk drives of the system.

TABLE 4-9 Disk Drives of the System

Feature	Description
Drive population	Each storage tray can have up to 14 disk drives.
Mixed drive size support	Mixing disk drive sizes is allowed but not recommended. If disk drives of different sizes are mixed in a storage tray, all the drives are treated as the smallest drive. For example, in a tray with a disk drive of 36 Gbytes and a drive of 146 Gbytes, only 36 Gbytes of the second disk drive is used. When the system is operating with mixed drive size support, the hot-spare must be the size of the largest disk drive in the tray.
Drive capacities	36 GB/15 KB, 73 GB/10 KB, 146 GB/10 KB
Online addition/removal of disks	Yes
Array hot-spare	Dedicated; that is, located in the array configuration and available as a spare to any virtual disk in the array configuration.
Dedicated hot-spare	Dedicated to the virtual disk; that is, located in the same array as the virtual disk.
Minimum number of disks	14
Maximum number of disks	448

For information on the current disk configuration, see [“Displaying Disk Information” on page 190](#).

Displaying Disk Information

You can display summary and detailed information about disks in your system.

To display information about disks:

1. Click Sun StorEdge 6920 Configuration Service > Physical Storage > Disks.

The Disk Summary page is displayed.

2. Select a disk to display detailed information.

The Disk Details page is displayed.

3. Go to Additional Information to see the number of replication sets and volumes associated with the selected disk.

Click an item to view its Summary page.

Managing External Storage

This section describes accessing user data or using storage that is available on a storage device that is external to the base cabinet. It contains the following subsections:

- [“About Using External Storage” on page 191](#)
- [“Planning to Use External Storage” on page 194](#)
- [“Displaying External Storage Information” on page 199](#)
- [“Creating a Legacy Volume” on page 200](#)
- [“Migrating Externally Stored Data” on page 201](#)
- [“Using External Storage as Raw Storage” on page 202](#)

About Using External Storage

You can access data and storage from a storage device that is not within the main base cabinet or an expansion cabinet. This enables you to:

- Migrate data from an external storage device to the system’s internal storage
- Manage and provide data services by using an external storage device’s logical unit number (LUN) as a volume (known as a legacy volume) on the system
- Add to the system’s storage capacity by using storage on an external storage device as raw storage

The following figure illustrates how a LUN on an external storage device can be used by the system:

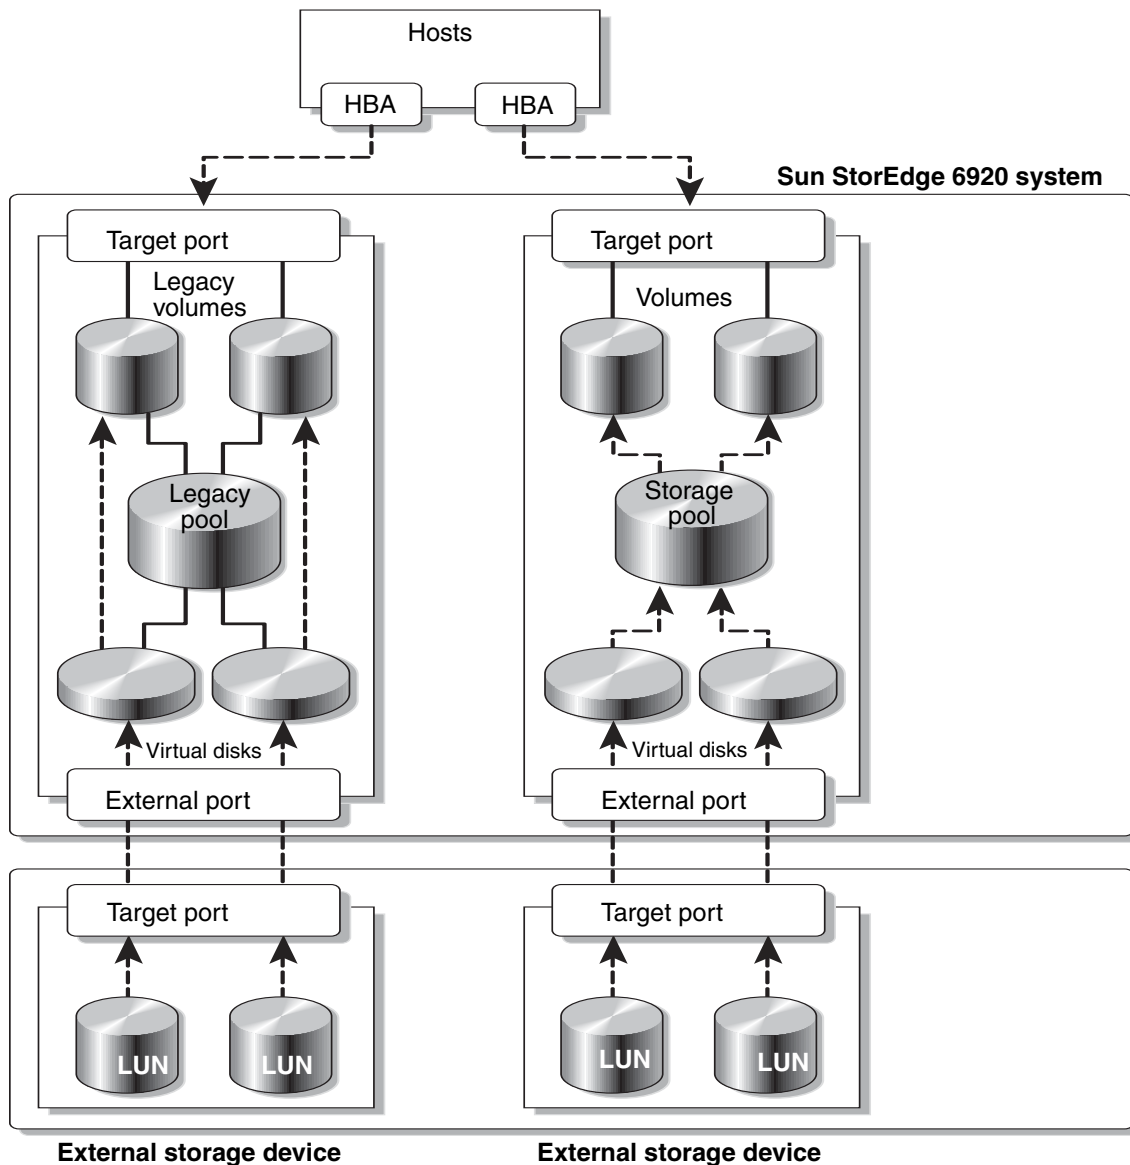


FIGURE 4-10 How a LUN and an External Storage Device Can Be Used by the System

You can connect external storage devices either directly to the system's service panel or indirectly through a single pair of redundant Fibre Channel (FC) switches. To ensure data reliability, create redundant paths through direct or FC switch connections.

Note: You cannot connect more than one set of redundant FC switches.

Currently supported external storage devices include:

- Sun StorEdge T3 array with 1 Gbyte cache
- Sun StorEdge 3510 array
- Sun StorEdge 3511 array
- Sun StorEdge 6120 array
- Sun StorEdge 6130 array
- Sun StorEdge 6320 system
- Sun StorEdge 9980 system
- Sun StorEdge 9970 system
- Sun StorEdge 9960 system
- Sun StorEdge 9910 system
- Sun StorEdge 3910/3960 systems (with Sun StorEdge T3 arrays with 1 Gbyte cache)
- Sun StorEdge 6910/6960 systems
- EMC CLARiiON CX 400 array
- EMC CLARiiON CX 700 array
- Hewlett Packard EVA 3000 virtual array

For an up-to-date list of supported external storage devices, see the *Sun StorEdge 6920 System Release Notes*.

For more information about how to cable an external storage device to the base cabinet, see the *Sun StorEdge 6920 System Getting Started Guide*.

Note: The system management software does not configure external storage devices. You must use the native configuration tools for the external storage device to make changes to external LUNs and lower-level components. For more information, see the documentation that ships with the external storage device.

When you connect an external storage device to the base cabinet, the system automatically discovers the storage device. The External Storage Summary page lists an external storage device's LUN as an uninitialized disk, which you can manage by using the Import External Storage wizard. When you have completed the wizard steps, the system perceives the external storage as a virtual disk, which you can use to do the following:

- Add to the system's storage capacity

You can designate storage on the external storage device as raw storage that is available for use by the system. Therefore, the system can use both of the following:

- Internal storage, which the arrays within the base cabinet provide

- External storage, which the external storage devices provide
- Create a legacy volume

A legacy volume consists of an entire external LUN on the external storage device. The system presents a legacy volume as a single local volume; however, the user data resides on the external storage device. You can use a legacy volume to do the following:

- Preserve the user data and manage the legacy volume as if it were any other volume on the system
- Preserve and migrate the user data from the external storage device to the internal storage in the system

After you create a legacy volume, the External Storage, Volume Summary, and Volume Details pages list the legacy volume as if it were any other volume. You can perform the following volume operations on a legacy volume:

- Map the legacy volume to an initiator
- Use the legacy volume to create a mirror component
- Add the legacy volume to an existing a mirror for resilvering (which destroys any existing data that resides on the legacy volume)
- Create a volume snapshot of the legacy volume and use the internal storage to hold the snapshot
- Roll back volume snapshot data to restore the legacy volume
- Move the legacy volume between legacy pools

Note: You cannot extend a legacy volume as you would any other volume on the system. Do not attempt to extend a legacy volume by using the external storage device's management software to increase the size of a LUN.

You can delete a legacy volume from the Volume Summary or Volume Details page. When you delete a legacy volume, it reverts to being an uninitialized LUN on the external storage device. The external storage device retains any data written to the legacy volume; if no data was written to the legacy volume, the LUN retains its original data completely.

Planning to Use External Storage

Before you can manage external storage, you must do the following:

- Use the external storage device's native configuration tools to ensure that the external storage device includes at least one logical unit number (LUN):
 - A LUN cannot be mapped to more than one port on the system.

- A LUN cannot appear on more than two controller ports.
- To support multipath failover, the LUN must appear on two of the external storage device's controller ports.
- For information about how to connect an external storage device to the system, see the *Sun StorEdge 6902 System Getting Started Guide*.
- If the external storage device is connected directly to the system, make sure that any LUNs mapped to the system are not accessible by other hosts.
- If the external storage device is connected to the system through redundant Fibre Channel (FC) switches, you must apply zoning practices. See the FC switch vendor's documentation for information about zoning. Consider the following when you plan to apply zoning practices:
 - Create zones so that each controller on the external storage device has an exclusive path to a single port on the system's service panel. You must ensure that a host is not mapped to the same LUN that is presented to the system.

The following figure shows a sample configuration in which the external storage device is connected to the system's service panel by means of a pair of FC switches. Each FC switch includes a zone with an exclusive path from a controller on the external storage device to a port on the system's service panel.

Note: A host can be attached to the switch, but it must not be mapped to the same LUN that is being presented to the system, and it cannot be in a zone that includes the system.

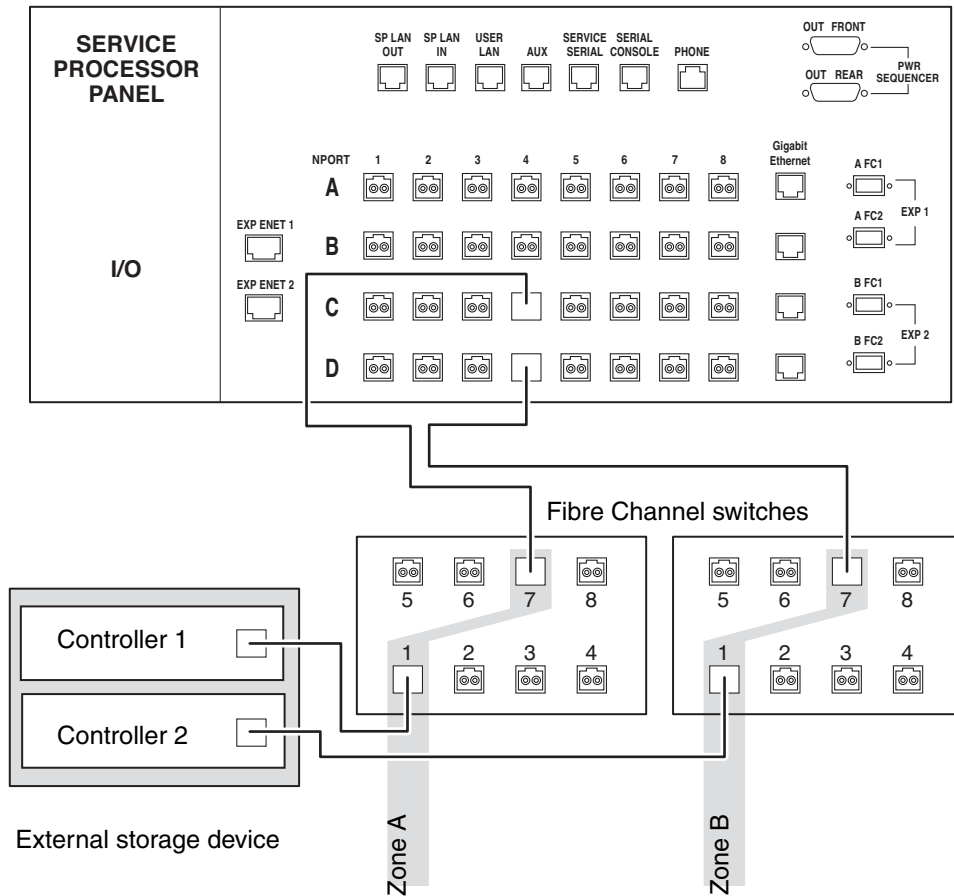


FIGURE 4-11 External Storage Device Connected to the System's Service Panel by FC Switches

- Create zones appropriately to filter out unwanted external storage devices so that the External Storage Summary page does not list an unmanageable number of LUNs.
- Initialize at least two virtual disks from the system's internal storage. If you create legacy volumes, the system requires two initialized virtual disks for administrative purposes. When you create a legacy volume, you cannot delete or uninitialized the last two remaining virtual disks on the system unless you delete the legacy volume first.
- Note that the system designates the entire LUN on the external storage device as a legacy volume. You cannot specify more than one LUN per legacy volume.

There are a number of decisions you must make before you run the Import External Storage wizard:

- Determine your user data preservation strategy:
 - If you want to preserve existing user data on a LUN on the external storage device, create a legacy volume. This operation ensures that the system does not overwrite existing user data by using it as raw storage.
 - If you do not need to preserve existing user data on a LUN on the external storage device, designate the storage as raw storage for use by the system. This operation overwrites the existing user data.
- If you are creating a legacy volume, choose a legacy volume name with a meaningful description for your environment.
- Determine the storage pool to which you want to add the legacy volume or raw storage:
 - If you create a legacy volume, the system lists the available legacy pools to which you can add the legacy volume. A legacy pool is defined by the Legacy profile, which ensures that you can manage the legacy volume and not allocate it as raw storage for new volumes. The wizard does not allow you to place a legacy volume in a pool other than a legacy pool.

The system provides a default legacy pool. Alternatively, you can create your own legacy pool, which must be defined by the preconfigured Legacy profile; you cannot create a custom Legacy profile.

Snapshot reserve space for a legacy volume must reside in a storage pool other than a legacy pool. You cannot allocate snapshot reserve space for a legacy volume from the pool in which it resides, that is, a legacy pool.
 - If you designate the external storage as raw storage, the system lists the available pools to which you can add the raw storage. The list includes all the storage pools on the system, other than legacy pools. Because the system cannot determine the attributes of a LUN on the external storage device, make sure that the LUN's attributes are consistent and compatible with the profile of the storage pool to which you plan to add the raw storage. For example, you do not want to add RAID-1 external storage to a pool defined by the RAID-5 profile.

After you have run the wizard, be aware of the following:

- The Additional Information table in the Volume Details page for a legacy volume always displays one virtual disk, which represents the LUN on the external storage device from which you created the legacy volume. The Management column of the Virtual Disk Summary page labels a LUN, which is managed by the system, as External. The following figure shows the relationship between LUNs on an external storage device and the legacy volumes on the system.

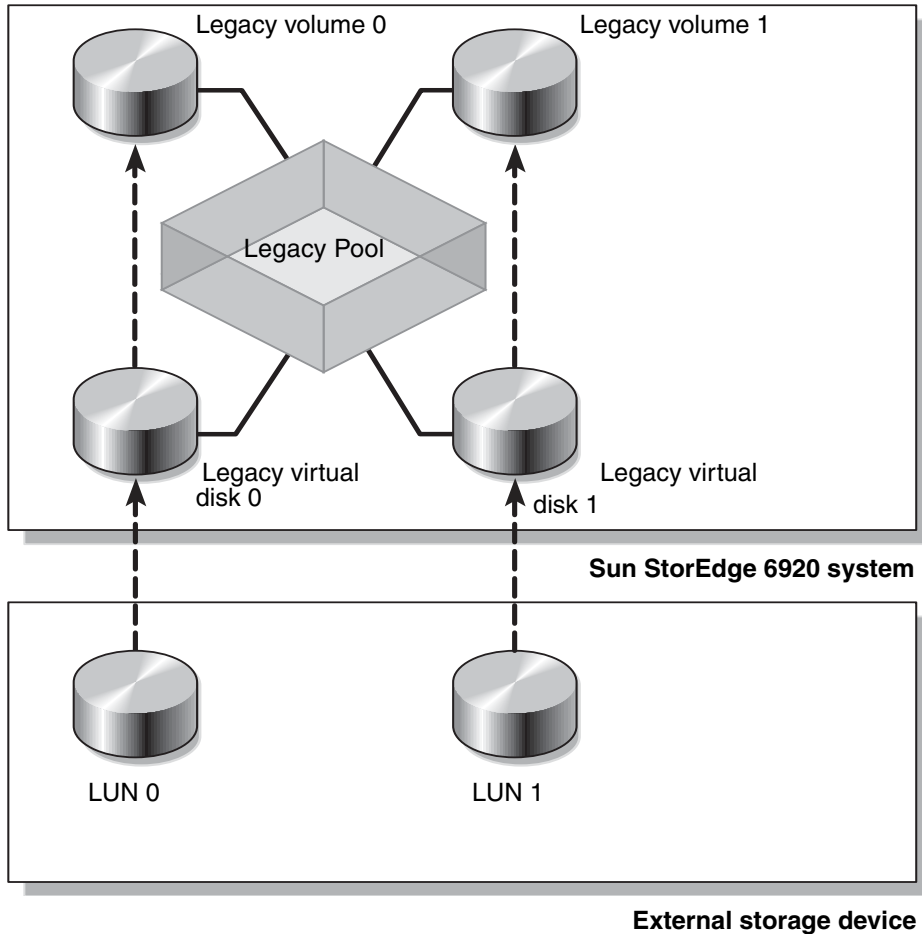


FIGURE 4-12 Relationship Between LUNs on an External Storage Device and Legacy Volumes on the System

- After you create a legacy volume, you have full access to the user data. You must map the legacy volume's WWN, which the system provides, to access the data.
- Make sure that you do not exceed the supported system limits when you connect to an external storage device:
 - The system supports 1024 volumes, which includes legacy volumes.
 - Each port can support 128 legacy volumes, as long as there are no initialized virtual disks active on the port.

- Unlike internal storage, the data paths to the external storage devices might not be redundant. Therefore, if the data path between the system and an external storage device becomes unavailable, you risk losing access to your data. A best practice is to provide redundant paths between the system and the external storage device to minimize this risk.

Note: The system does not provide diagnostic support for problems incurred as a result of connections to external storage devices.

If a legacy volume or external storage is unavailable, do the following:

- Verify that your SAN is configured properly.
- Verify that you have zoned your fabric appropriately.
- Verify that the external storage device is cabled to the system properly.
- Contact Sun Professional Services.

For more information about SANs, FC switches, and zoning, see the Related Documentation section of the *Sun StorEdge 6920 System Getting Started Guide*.

Displaying External Storage Information

So that the External Storage Summary page does not list an unmanageable number of logical unit numbers (LUNs), apply zoning or network masking best practices to filter out unwanted external storage devices. See the vendor's documentation for information about zoning.

To display information about LUNs that are available to the system:

1. Click Sun StorEdge 6920 Configuration Service > External Storage.

The External Storage Summary page is displayed.

If the Pool and Legacy Volume columns show no information, the LUNs appear as uninitialized disks and must be managed by means of the Import External Storage wizard before you can access the external storage or user data.

If a legacy volume or raw storage has been created from a LUN, the system lists the pool in which it has been placed and, in the case of a legacy volume, its name in the Legacy Volume column.

2. In the Legacy column, click the name of a legacy volume.

The Details page displays information about the selected legacy volume as if it were any other volume on the system.

Note: If you connect or disconnect an external storage device, the system detects the change and the External Storage Summary page displays the updated information. If the configuration has changed and the External Storage Summary page does not display what you expect, click the Rescan Devices button. The software performs a scan of all storage and host ports for changes in device configuration. The software might take 10 to 30 seconds to update the External Storage Summary page.

Creating a Legacy Volume

You must consider a number of factors and make a number of decisions before using external storage as raw storage. For information, see “Planning to Use External Storage” on page 194.

To preserve the user data on a LUN on the external storage device, and consequently create a legacy volume:

1. On the external storage device that you plan to manage, disconnect or unmap any host that has access to the LUN.

If the system shares a LUN on the external storage device with another host, and you do not disconnect that host, the user data might become corrupted.

2. Make sure the system includes at least two initialized internal virtual disks, which the system requires to maintain the metadata for the legacy volumes.
3. Click Sun StorEdge 6920 Configuration Service > External Storage.

The External Storage Summary page is displayed.

4. Select a LUN or LUNs on the external storage device that you want to access.

If the system lists a LUN that is currently managed by the system and perceived as a virtual disk, the entry does not have a check box.

5. Click Import.

The Import External Storage wizard is displayed.

6. Follow the steps in the wizard, making sure to choose the option to preserve the existing data on the external storage device.

Migrating Externally Stored Data

You must consider a number of factors and make a number of decisions before migrating user data from an external storage device. For information, see “Planning to Use External Storage” on page 194.

You can combine the management of external storage with the system’s data mirroring feature to migrate user data from an external storage device onto the system.

To migrate user data:

1. Disconnect or unmap any host that has access to the LUN on the external storage device that you plan to manage.

If the system shares a LUN on the external storage device with another host, the user data might be corrupted.

2. Make sure the system includes at least two initialized internal virtual disks, which the system requires to maintain the metadata for the legacy volumes.
3. Click Sun StorEdge 6920 Configuration Service > External Storage.

The External Storage Summary page is displayed.

4. Select a LUN or LUNs on the external storage device that you want to access.

If the system lists a LUN that is currently managed by the system and perceived as a virtual disk, the entry does not have a check box.

5. Click Import.

The Import External Storage wizard is displayed.

6. Follow the directions in the wizard, making sure to choose the option to preserve the existing data on the external storage device.
7. Map the newly created legacy volume to a host, if applicable.

Note: When you map a host to the legacy volume, a host application can access user data immediately and continue to access user data while you perform the following steps to complete the data migration.

8. Use the data mirroring feature to copy the user data on the legacy volume to a separate mirror component, which comes from an available storage pool on the system.

Because the smallest volume restricts the size of a mirror, make sure you know the size of the legacy volume. The mirror components must be the same size.

You might want to create several new components as part of the mirror so that there are multiple independent copies of the user data in internal storage to use for protection, backup, and the like.

You can view Jobs page to monitor the status resilvering process. Refer to “Displaying Configuration Job Information” on page 230 for information about a job’s status.

When the resilvering process is complete, the user data is available from the legacy volume and also from the new mirror component, which is part of the system’s internal storage. For more information about how to use the system’s data mirroring feature, see “Planning Mirror Components” on page 95 and “Mirroring an Existing Volume” on page 98.

9. Break the legacy volume component from the mirror, as described in “Breaking a Mirror Component” on page 101.

You will have two independent copies of the user data:

- The legacy volume, which is external storage
- The mirror component of the legacy volume, which is internal storage

Note: The application continues to access the user data from the mapped host. Because you broke the legacy volume component from the mirror, the system finds the data exclusively on the mirror component in internal storage.

10. Remove the external storage device from the base cabinet, if applicable.

Using External Storage as Raw Storage

You must consider a number of factors and make a number of decisions before using external storage as raw storage. For information, see “Planning to Use External Storage” on page 194.

Note: If you designate the external storage as raw storage, the system will not preserve pre-existing user data.

To designate external storage as raw storage:

1. Determine the attributes of the LUNs on the external storage devices to ensure that they are compatible with profiles of the pools in which you plan to place the raw storage.
2. Click Sun StorEdge 6920 Configuration Service > External Storage.
The External Storage Summary page is displayed.
3. Select the LUN or LUNs that you want to manage.

4. Click Import.

The Import External Storage wizard is displayed.

5. Follow the steps in the wizard, making sure to choose the option to add the storage to a pool as raw storage.

Monitoring the System

This chapter describes system monitoring tasks. It contains the following sections:

- [“Setting Diagnostic Viewing Preferences” on page 206](#)
- [“Monitoring Alarms and Events” on page 207](#)
- [“Monitoring Devices” on page 213](#)
- [“Monitoring Topologies” on page 225](#)
- [“Managing Storage Configuration Jobs” on page 228](#)
- [“Managing Storage Monitoring Jobs” on page 232](#)

Setting Diagnostic Viewing Preferences

You can set display preferences for how you want to view diagnostic topologies and asset reports.

To set preferences:

1. Click Sun Storage Automated Diagnostic Environment.
2. In the banner, click Preferences.

The Edit Preferences window is displayed.

3. Configure the preference options listed in the following table. The following table describes the buttons and fields on the Edit Preferences window.

TABLE 5-1 Edit Preferences Window

Field	Description
<i>Topology Display</i>	
Default Topology View Type	Determines whether the default topology view is Graphical, Tabular, or Interactive Applet.
Preferred Object Label	Determines whether system objects are identified by Device Name, IP Name/Address, or Vendor/Model.
Preferred Object Tool Tip	Determines whether system object tooltips are identified by Device Name, IP Name/Address, or Vendor/Model.
Object Scaling Behavior	Determines how objects are grouped during topology scaling operations. Options are: <ul style="list-style-type: none">• Disable Grouping – Objects are not grouped during scaling operations.• Enable Grouping – Objects are grouped during scaling operations.• Enable Nested Grouping – Objects are grouped in a nested relationship during scaling operations.
<i>Asset Report Display</i>	
Default Asset Details View Type	Determines whether the default Device Details page view is in summary or detailed form.

4. Click Save.
A confirmation message is displayed.
5. Click Close.

Monitoring Alarms and Events

This section describes alarms and events. It contains the following subsections:

- [“About Alarms and Events” on page 207](#)
- [“About Alarm Management” on page 207](#)
- [“Displaying Alarm Information” on page 208](#)
- [“Acknowledging Alarms” on page 210](#)
- [“Reopening Acknowledged Alarms” on page 210](#)
- [“Deleting Alarms” on page 211](#)
- [“Generating Alarm Statistics Reports” on page 211](#)
- [“Displaying Event Information” on page 212](#)

About Alarms and Events

Events are generated to signify a health transition in a monitored device or device component. Events that require action are classified as alarms.

There are four event severity levels:

- **Down** – Identifies a device or component as not functioning and in need of immediate service
- **Critical** – Identifies a device or component in which a significant error condition is detected that requires immediate service
- **Major** – Identifies a device or component in which a major error condition is detected and service may be required
- **Minor** – Identifies a device or component in which a minor error condition is detected or an event of significance is detected

About Alarm Management

Alarms are not automatically deleted from the Alarms page. You must manually delete any alarm that you want to remove from the Alarms page. If you do not remove alarms that are no longer current, new, more recent minor alarms for a component will not be displayed on the Alarms page in place of previous, more major alarms for the same component.

When you are running the Enterprise Edition in conjunction with the System Edition and you have alarms that are passed from the System Edition to the Enterprise Edition, you must manually delete these alarms from the browser interfaces of both editions. If you do not do so, new lower-severity alarms for any component being monitored by the Enterprise Edition software will not display in place of previous higher-severity alarms for the same component.

Displaying Alarm Information

This section describes how to display alarm information. It contains the following subsections:

- [“Displaying an Alarm Summary” on page 208](#)
- [“Displaying Alarm Details” on page 208](#)
- [“Displaying Aggregated Events” on page 208](#)
- [“Displaying Alarms for a Specific Device” on page 209](#)
- [“Displaying an Alarm Statistics Report” on page 209](#)

Displaying an Alarm Summary

To display alarm summary information, click File System Manager > Alarms.

The Alarm Summary page is displayed.

Displaying Alarm Details

To display detailed information about an alarm:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.
The Alarm Summary page is displayed.
2. Click Details for the alarm for which you want to display detailed information.
The Alarm Details page is displayed for the selected alarm.

Displaying Aggregated Events

To display aggregated events for an alarm:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.
The Alarm Summary page is displayed.
2. Click Details for the alarm for which you want to display detailed information.
The Alarm Details page is displayed for the selected alarm.
3. Click View Aggregated Events.
A list of all events associated with the displayed alarm is displayed.

Displaying Alarms for a Specific Device

To display detailed alarm information for a specific device:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.
The Alarm Summary page is displayed.
2. In the Device column of the Alarm Summary table, click the name of the device for which you want to display alarms.
The Device Details page is displayed.

Displaying an Alarm Statistics Report

Alarm statistics provide event and alarm trends over days of the week, weeks, and months.

To display an alarm statistics report:

1. Click Sun Storage Automated Diagnostic Environment > Alarms > Statistics.
The Alarms Statistics page is displayed.
2. In the Actions column, click View for one of the following alarm statistic report types:
 - By device type and month
 - By device type and week
 - By device type and day of weekThe selected alarm statistics report is displayed.

Acknowledging Alarms

When an alarm is generated, it remains open in the Alarm Summary page until you acknowledge it. Acknowledged alarms remain in the Alarm Summary page until you delete them.

To acknowledge one or more alarms:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.

The Alarm Summary page is displayed.

2. Select the check box for each alarm you want to acknowledge.

3. Click Acknowledge.

The Acknowledge Alarms confirmation window is displayed.

4. Enter an identifying name to be associated with this action.

5. Click Acknowledge.

The Alarm Summary page is redisplayed, and the state of the acknowledged alarms is displayed as Acknowledged.

Reopening Acknowledged Alarms

To reopen one or more previously acknowledged alarms:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.

The Alarm Summary page is displayed.

2. Select the check box for each acknowledged alarm you want to reopen.

3. Click Reopen.

The Reopen Alarms confirmation window is displayed.

4. Enter an identifying name to be associated with this action.

5. Click Reopen.

The Alarm Summary page is redisplayed, and the state of the reopened alarms is displayed as Open.

Deleting Alarms

When you delete an open or acknowledged alarm, it is permanently removed from the Alarm Summary page.

To delete one or more alarms:

1. Click Sun Storage Automated Diagnostic Environment > Alarms.
The Alarm Summary page is displayed.
2. Select the check box for each alarm you want to delete.
3. Click Delete.
The Delete Alarms confirmation window is displayed.
4. Click OK.
The Alarm Summary page is redisplayed without the deleted alarms.

Generating Alarm Statistics Reports

Alarm statistics provide event and alarm trends over past days of the week, weeks, and months.

To generate updated alarm statistics reports:

1. Click Sun Storage Automated Diagnostic Environment > Alarms > Statistics.
The Alarm Statistics page is displayed.
2. Click Generate All Reports.
When the reports are complete, the following confirmation message is displayed:
`Alarm Statistics Reports have been generated.`

Displaying Event Information

This section describes how to display event information. It contains the following subsections:

- [“Displaying the Event Log” on page 212](#)
- [“Displaying Event Details” on page 212](#)

Displaying the Event Log

To gather additional information about an alarm, you can display the event log to view the underlying events on which the alarm is based.

To display event log information, click File System Manager > Administration > Event Log.

The Events page displays a summary of all events in the system event log.

Displaying Event Details

To display event details:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Event Log.

The Events page displays a summary of all events in the system event log.

2. Click Details in the row that corresponds to the event for which you want detailed information.

The Event Details page is displayed for the selected event.

Monitoring Devices

This section describes how to monitor and manage devices. It contains the following subsections:

- [“Displaying Device Information” on page 213](#)
 - [“Displaying Performance Data” on page 216](#)
 - [“Adding Notes for a Rack” on page 218](#)
 - [“Adding Notes for a Device” on page 218](#)
 - [“Generating New Inventory” on page 219](#)
 - [“Changing the Monitoring Status for a Device” on page 219](#)
 - [“Setting Up Volume Verify” on page 220](#)
 - [“Installing and Uninstalling Revision Patches” on page 221](#)
-

Displaying Device Information

This section describes how to display device information. It contains the following subsections:

- [“Displaying a Rack Inventory Summary” on page 213](#)
- [“Displaying Rack Components” on page 214](#)
- [“Displaying Device Details” on page 214](#)
- [“Displaying Device FRUs, Components, or Logical Components of a Specific Type” on page 214](#)
- [“Displaying Detailed Information for a Specific Device FRU, Component, or Logical Component” on page 215](#)
- [“Displaying System Parameters for a Device” on page 216](#)

Displaying a Rack Inventory Summary

To display summary information about all system racks, click File System Manager > Inventory.

The Rack Inventory page displays summary information for each system device by rack.

Displaying Rack Components

To display a detailed list of all components in all system racks:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. From the Actions drop-down menu, select View Rack Components.
The Device Details page displays a detailed list of all field-replaceable units (FRUs), components, and logical components for all racks.

Displaying Device Details

To display detailed information for a specific device:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. Click the name of the rack-based device for which you want detailed information.
The Device Details page for the selected device is displayed.
3. (Optional) From the View drop-down menu, select Detailed.
The Device Details page for the selected device is redisplayed with detailed information.

Displaying Device FRUs, Components, or Logical Componentsof a Specific Type

To display a list of device field-replaceable units (FRUs), components, or logical components of a specific type:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. Click the name of the rack-based device for which you want detailed information about a specific component.
The Device Details page for the selected device is displayed.

3. Optionally, from the View drop-down menu, select Detailed.
The Device Details page for the selected device is redisplayed with detailed information.
4. Scroll down the Device Details page until the FRU Summary, Component Summary, Logical Component Summary, or Volume Summary for which you want information is listed, and click the name of item for which you want a list.
A Summary page for the selected type is displayed.

Displaying Detailed Information for a Specific Device FRU, Component, or Logical Component

To display information for a specific device field-replaceable unit (FRU), component, or logical component:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. Click the name of the rack-based device for which you want detailed information about a specific component.
The Device Details page for the selected device is displayed.
3. Optionally, from the View drop-down menu, select Detailed.
The Device Details page for the selected device is redisplayed with detailed information.
4. Scroll down the Device Details page until the FRU Summary, Component Summary, Logical Component Summary, or Volume Summary for which you want information is listed, and click the name of the item for which you want a list.
A summary page for the selected type is displayed.
5. Click the name of the specific FRU, component, logical component, or volume for which you want detailed information.
A Details page for the selected item is displayed.

Displaying System Parameters for a Device

To display system parameters for a device:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each device by rack.
2. Click the name of the rack-based device for which you want to display system parameters.
The Device Details page for the selected device is displayed.
3. From the Actions drop-down menu, select View System Parameters.
System parameters for the selected device are displayed.
Note: System parameters are not available for all devices.

Displaying Performance Data

This section describes how to display performance data. It contains the following subsections:

- [“Displaying Rack Performance Data” on page 216](#)
- [“Displaying Performance Data for a DSP or Array” on page 217](#)
- [“Displaying DSP Utilization or Replication Queue Depth” on page 217](#)

Displaying Rack Performance Data

To display performance data for the Data Services Platform (DSP) and all arrays in a rack:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. From the Actions drop-down menu, select View Performance Data.
The Rack Performance page displays performance data for the DSP and for each array in the rack.
3. (Optional) Click the Advanced Filter icon to filter the performance data by specific DSP ports or selected arrays.

4. (Optional) To change the way performance data is displayed, do any of the following ways:
 - From the Report Type drop-down menu, select Current Value to display the current performance data or select Hourly Average to display performance data averaged over the last 24 hours.
 - From the View drop-down menu, select I/Os per Second to display the number of Read and Write input/output (I/O) operations per second or select MB per Second to display the size in megabytes (MB) of read and write operations.
 - From the Layout drop-down menu, select Tabular to display the performance data in a table or select Graphical to display the performance data in a graph.

Displaying Performance Data for a DSP or Array

To display performance data for a specific Data Services Platform (DSP) or array:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. Click the name of the DSP or array for which you want to display performance data.
The Device Details page for the selected device is displayed.
3. From the Actions drop-down menu, select View Performance Data.
The Device Performance page displays performance data for the selected DSP or array.

Displaying DSP Utilization or Replication Queue Depth

Note: You can display replication queue depth only if remote replication is enabled. For more information, see “Using Data Replication” on page 116.

To display Data Services Platform (DSP) utilization or replication queue depth:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.

2. Click the name of the DSP for which you want to display more information.
The Device Details page for the DSP is displayed.
3. From the Actions drop-down menu, select View Performance Data.
The Device Performance page displays performance data for the selected DSP.
4. From the View drop-down menu, select one of the following:
 - Select Utilization to display the read and write utilization percentage of each DSP port.
 - Select Replication Queue Depth to display the replication queue depth for each volume, including current queue depth and queue capacity.

Adding Notes for a Rack

You can add notes to information about a rack in the system. This enables you to keep a record of configuration changes or service tasks for the rack.

To add notes to be stored with rack information:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. Go to the Notes section of the page.
3. Type your notes in the Notes text box, and click Save.

Note: To reset the contents of the Notes text box to the last saved notes, click Reset.

Adding Notes for a Device

You can add notes to information about a device in the system. This enables you to keep a record of configuration changes or service tasks for the device.

To add notes to be stored with device information:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.

2. Click the name of the device for which you want to add a note.
The Device Details page for the selected device is displayed.
3. Go to the Notes section of the page.
4. Type your notes in the Notes text box and click Save.

Note: To reset the contents of the Notes text box to the last saved notes, click Reset.

Generating New Inventory

Generate Inventory is used to update the list of monitored devices when devices are added or removed from a rack or when an expansion rack is added or removed. Both the Inventory and Topology pages will reflect the changes after the inventory is generated.

To generate new inventory:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page is displayed.
2. From the Actions drop-down menu, select Generate New Inventory.
The Generate Inventory page is displayed.
3. Click Generate New Inventory.
A table of new inventory changes is displayed.
4. Click Save.
The inventory changes are saved to the inventory database.

Changing the Monitoring Status for a Device

You can quickly change the monitoring status for a device using the Update Monitoring and Setup Information.

To update the monitoring status information for a device:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each device by rack.
2. Click the name of the rack-based device for which you want to update monitoring and setup information.
Details for the selected device are displayed.
3. From the Actions drop-down menu, select Update Monitoring and Setup Information.
The Monitoring and Setup Information pop-up window is displayed for the device.
4. To disable monitoring of the device, deselect the Monitor check box or, to enable monitoring of the device, select the Monitor check box.
5. Click OK.
The inventory is updated to reflect specified change.

Setting Up Volume Verify

You can set up the Volume Verify function to run manual parity checks on array volumes.

Note: Disk scrubber must be turned off for Volume Verify to be enabled. For instructions, see “Modifying an Array” on page 183.

To set up the volume verify function to run on all arrays:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each device by rack.
2. From the Actions drop-down menu, select Set Up Volume Verify.
The Volume Verify Setup page is displayed.
3. For the Activate Volume Verify option, select Yes.
4. To enable Volume Verify to automatically fix errors it encounters while running, select Yes for the “Run with fix option” field.
Note: This option might mistakenly fix parity on bad blocks. Contact your storage system representative for more information.

5. From the “Run on same array at most every X days” list, select the maximum frequency, in days, with which you want Volume Verify to run on any one array.
6. From the day-of-the-week lists, choose a time period from at least one day at which that you want Volume Verify to run.
Note: This parameter determines the scheduled start time of the volume verify function. The function may continue beyond the selected time period.
7. From the Execution Rate list, select the rate at which you want Volume Verify to run.
8. Enter the default array password, if configured, for the system.
9. Click OK.

Installing and Uninstalling Revision Patches

This section describes how to install and uninstall revision patches. It contains the following subsections:

- [“About Installing and Uninstalling Revision Patches” on page 221](#)
- [“Generating a Revision Maintenance Upgrade Report” on page 222](#)
- [“Installing Revision Patches” on page 222](#)
- [“Generating a Revision Maintenance Backout Report” on page 223](#)
- [“Backing Out Revision Patches” on page 224](#)

About Installing and Uninstalling Revision Patches

The revision maintenance function enables you to upgrade patches and back out installed patches.

- **Revision Upgrade** – The revision upgrade function enables you to generate a report of available patches for system devices and to install patches for one or more devices. Keep the following guidelines in mind when installing patches:
 - You can choose to install only a subset of available device patches.

- When you finish installing one or more patches, it is a good idea to regenerate the revision maintenance upgrade report to determine whether there are additional patches available.
- Revision Backout – The revision backout function enables you to generate a list of the last installed patch for each device and to uninstall one or more device patches. Keep the following guidelines in mind when backing out patches:
 - You can back out only the last installed patch for each device.
 - Only patches installed done through the revision maintenance upgrade function can be backed out. If a patch is installed as part of a field-replaceable unit (FRU) replacement or addition done through the Service Advisor, the patch cannot be backed out.

Generating a Revision Maintenance Upgrade Report

You can generate a report listing the required patches for all devices in the storage system.

To generate a revision maintenance upgrade report:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Rack Inventory page displays summary information for each system device by rack.
2. From the Actions drop-down menu, select Revision Maintenance Upgrade.
The Revision Maintenance Upgrade page is displayed. If an upgrade report had been previously generated, the required patches for system devices are shown.
3. To generate an updated upgrade report, click Generate Upgrade Report.
The Revision Maintenance Upgrade page displays an updated list of required patches for system devices. You can monitor the progress of the report generation using the Jobs function.

Installing Revision Patches

You can install one or more patches for devices in the storage system.

To install one or more revision patches:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.

The Rack Inventory page displays summary information for each system device by rack.

2. From the Actions drop-down menu, select Revision Maintenance Upgrade.

The Revision Maintenance Upgrade page displays an updated list of required patches for system devices.

3. Click Generate Upgrade Report.

The Revision Maintenance Upgrade page displays available patches for system devices.

4. (Optional) To display the volumes that will be affected when a patch is installed on an array, click View Affected Volumes.

The Affected Volumes page displays the volumes that will be affected for each disk patch.

5. Select the check box for each device patch you want to install, and click Apply.

The selected patches are installed. You can monitor the progress of the patch installation using the Jobs function.

Note: After installing patches, you should regenerate the Revision Maintenance Upgrade Report to determine whether any additional patches are available.

Generating a Revision Maintenance Backout Report

You can generate a report listing the installed patches for all devices in the storage system.

To generate a revision maintenance backout report:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.

The Rack Inventory page displays summary information for each system device by rack.

2. From the Actions drop-down menu, select Revision Maintenance Backout.

The Revision Maintenance Backout page is displayed. If a backout report has been previously generated, the installed patches for system devices are shown.

3. To generate an updated backout report, click Generate Backout Report.

The Revision Maintenance Backout page displays an updated list of the last installed patches for system devices. You can monitor the progress of the report generation using the Jobs function.

Note: You can uninstall only the last installed patch for each device, and you can uninstall only patches that were installed using the revision maintenance upgrade function.

Backing Out Revision Patches

You can back out all installed revision patches on all system devices.

Note: You cannot back out individual patches on individual devices.

To back out all revision patches:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.

The Rack Inventory page displays summary information for each system device by rack.

2. From the Actions drop-down menu, select Revision Maintenance Backout.

The Revision Maintenance Backout page is displayed.

3. Click Generate Backout Report.

The Revision Maintenance Backout page displays an updated list of installed patches by system device.

4. (Optional) To display the volumes that will be affected when a patch is backed out of an array, click View Affected Volumes.

The Affected Volumes page displays the volumes that will be affected for each disk patch to be backed out.

5. Select the check box to the left of each device for which you want to back out the last installed patch.

Note: You can uninstall only the last installed patch for each device.

6. Click Backout All Patches.

All revision patches on all devices are uninstalled. You can monitor the progress of the backout function using the Jobs function.

Monitoring Topologies

This section describes how to monitor and manage devices in topology views. It contains the following subsections:

- [“Displaying the System Topology” on page 225](#)
- [“Changing Topology Views” on page 225](#)
- [“Archiving a Topology” on page 226](#)
- [“Displaying Archived Topologies” on page 226](#)
- [“Displaying an Archived Topology Snapshot” on page 226](#)
- [“Comparing Two Archived Topology Snapshots” on page 227](#)
- [“Deleting an Archived Topology Snapshot” on page 227](#)

Displaying the System Topology

To display the system inventory in a topology view, click File System Manager > Topology.

The system configuration is displayed in the default topological view set in Preferences.

Changing Topology Views

To change the topology view from graphical or tabular:

1. Click File System Manager > Topology.

The system configuration is displayed in the default topological view set in Preferences.

2. From the View drop-down menu, select the desired view type.

The Topology page is redisplayed with the new topology view.

Archiving a Topology

To archive the current topology:

1. Click File System Manager > Topology.

The system configuration is displayed in the default topological view set in Preferences.

2. From the Actions drop-down menu, select Archive Current Topology.

A confirmation is displayed when the topology is successfully archived.

Displaying Archived Topologies

To display a list of archived topology views:

1. Click File System Manager > Topology.

The system configuration is displayed in the default topological view set in Preferences.

2. From the Actions drop-down menu, select View Archived Topologies.

A list of archived topologies is displayed.

Displaying an Archived Topology Snapshot

To display an archived topology snapshot:

1. Click File System Manager > Topology.

The system configuration is displayed in the default topological view set in Preferences.

2. From the Actions drop-down menu, select View Archived Topologies.

A list of archived topologies is displayed.

3. Select the check box of the archived topology you want to display, and click View.
The selected topology is displayed.

Comparing Two Archived Topology Snapshots

To compare two archived topology snapshots:

1. Click File System Manager > Topology.
The system configuration is displayed in the default topological view set in Preferences.
2. From the Actions drop-down menu, select View Archived Topologies.
A list of archived topologies is displayed.
3. Select the check boxes of the archived topologies you want to compare, and click Compare.
The Compare Archived Topologies page lists differences between the selected archived topologies.

Deleting an Archived Topology Snapshot

To delete an archived topology snapshot:

1. Click File System Manager > Topology.
The system configuration is displayed in the default topological view set in Preferences.
2. From the Actions drop-down menu, select View Archived Topologies.
A list of archived topologies is displayed.
3. Select the check box of the archived topology you want to delete, and click Delete.
The Archived Topologies page is redisplayed, and the selected topology is no longer listed.

Managing Storage Configuration Jobs

This section describes storage configuration jobs. It contains the following subsections:

- [“About Configuration Jobs” on page 228](#)
- [“Reference: Configuration Job Triggers” on page 228](#)
- [“Displaying Configuration Job Information” on page 230](#)
- [“Canceling and Removing Configuration Jobs” on page 231](#)

About Configuration Jobs

When you perform an operation on an object, there are cases in which the system does not process the operation immediately, but creates a job. The system creates a configuration job if you perform the following:

- Multiple operations, such as deleting a number of volumes at one time. Because operations on more than one object can affect performance, the system creates a job that completes the operation while you make other selections.
- Operations that by nature run for extended periods of time, such as resilvering a mirror.

The system lists these configuration jobs on the Job Summary page, which enables you to follow their progress and keep a history of their operations. While a job is in progress, it is reported on the Current Job Summary page. When a job is completed, the Historical Job Summary page lists the job.

Note: The completion of jobs can be affected by the browser’s session time. If the browser session time expires, or if you log out of the Java Web Console, the job terminates and the system reports it as a historical job with a status of Canceled.

Reference: Configuration Job Triggers

If you perform multiple operations, each operation has a threshold number that indicates how many selected objects cause the system to create a job.

The following table lists the operations that create a configuration job when the number of objects you select meets or exceeds the operation's threshold number.

TABLE 5-2 Operations That Can Create a Configuration Job Based on the Number of Objects Selected

Job Name	Operation	Origination of the Operation	Threshold Number
Volume Delete	Deleting a volume	Volume Summary	2
Volume Snapshot	Creating a volume snapshot	Volume Summary	2
Volume Map	Mapping a volume to an initiator	Volume Summary Volume Detail Create Volume wizard Create Snapshot wizard	2
Initiator Map	Mapping an initiator to a volume	Initiator Summary Initiator Details Create Initiator wizard	2
Mapping Delete	Deleting the mapping between an initiator and a volume	Initiator Details-Mapping Summary Volume Details-Mapping Summary	2
Pool Delete	Deleting a storage pool	Storage Pool Summary	2
Initiator Delete	Deleting an initiator	Initiator Summary	2
Profile Delete	Deleting a storage profile	Storage Profile Summary	2
Job Cancel	Canceling a job in progress	Current Job Summary	50
Job Remove	Removing a completed or canceled job	Historical Job Summary	50
Snapshot Delete	Deleting a snapshot	Snapshot Summary	2
Snapshot Resnap	Resnapping a volume snapshot	Snapshot Summary	2
Replication Set Delete	Deleting a replication set	Replication Set Summary	2

The following operations always create a configuration job because they might run for extended periods of time.

TABLE 5-3 Operations That Always Create a Configuration Job Due to Running Extended Periods of Time

Job Name	Operation	Origination of the Operation
Virtual Disk Delete	Deleting a virtual disk	Virtual Disk Summary
Resilvering Mirror Components	Creating a mirrored volume	Volume Summary New Volume wizard
Resilvering Mirror Components	Mirroring an existing volume	Volume Details New Mirrored Volume wizard
Resilvering Mirror Components	Adding a component to a mirrored volume	Mirrored Volume Details Add to Mirror wizard
Resilvering Mirror Component	Rejoining split components of a mirrored volume	Mirrored Volume Details Mirror Component Details
Resilvering Mirror Components	Rejoining split components fo a mirror	Mirrored Volume Details Mirror Component Details
Rollback Volume	Rolling a volume back to a snapshot	Snapshot Details
Virtual Disk Create	Adding storage to a storage pool	Storage Pool Summary Storage Pool Details Wizards with the option to add storage
Replication Synchronize	Synchronizing data within a replication set	Replication Set Details Consistency Group Details

Displaying Configuration Job Information

To display configuration job information:

1. Click Sun StorEdge 6920 Configuration Service > Jobs.
2. Click one of the following tabs, depending on the type of job:
 - Current
 - Historical

For more information about a specific job, click on its ID.

Canceling and Removing Configuration Jobs

To cancel a current configuration job:

1. Click Sun StorEdge 6920 Configuration Service > Jobs > Current.
The Current Job Summary page is displayed.
2. Select the job you want to cancel.
3. Click the Cancel Job button.

You can also cancel a job from the Current Job Details page.

Note: If you are planning to cancel a volume rollback job, be aware that there might be significant impact to your system. Refer to “About Snapshot Rollbacks” on page 105 for more information about how to manage volume rollbacks.

To remove a historical job:

1. Click Sun StorEdge 6920 Configuration Service > Jobs > Historical.
The Historical Job Summary page is displayed.
2. Select the job you want to remove.
3. Click the Remove button.

You can also remove a job from the Historical Job Details page.

Managing Storage Monitoring Jobs

This section describes storage monitoring and diagnostic jobs. It contains the following subsections:

- [“About Monitoring Jobs” on page 232](#)
- [“Displaying Monitoring Job Information” on page 232](#)
- [“Stopping a Monitoring Job” on page 234](#)
- [“Archiving Monitoring Jobs” on page 234](#)
- [“Deleting Monitoring Jobs” on page 234](#)
- [“Deleting Archived Monitoring Jobs” on page 235](#)

About Monitoring Jobs

All system processes that you initiate or schedule are tracked as system jobs. Running and completed jobs can be displayed in a summary list, and specific jobs can be displayed individually in detail.

Jobs can be archived for future reference.

Displaying Monitoring Job Information

This section describes how to monitor and manage devices in topology views. It contains the following subsections:

- [“Displaying a Summary of Monitoring Jobs” on page 232](#)
- [“Displaying the Details of a Monitoring Job” on page 233](#)
- [“Displaying a Summary of Archived Monitoring Jobs” on page 233](#)
- [“Displaying the Details of an Archived Monitoring Job” on page 233](#)

Displaying a Summary of Monitoring Jobs

To display a summary of pending and completed jobs, click File System Manager > Jobs.

The Job Summary page is displayed.

Displaying the Details of a Monitoring Job

To display the details of a completed monitoring job:

1. Click File System Manager > Jobs.
The Job Summary page is displayed.
2. Click Details in the row that corresponds to the job for which you want to display details.
The Job Details page is displayed.

Displaying a Summary of Archived Monitoring Jobs

To display a list of archived monitoring jobs, click File System Manager > Jobs > Archived Jobs.

The Archived Jobs page is displayed.

Displaying the Details of an Archived Monitoring Job

To display the details of an archived monitoring job:

1. Click File System Manager > Jobs > Archived Jobs.
The Archived Jobs page is displayed.
2. Click Details in the row that corresponds to the archived job for which you want to display details.
The selected Job Details page is displayed.

Stopping a Monitoring Job

To stop a monitoring job that is in progress:

1. Click File System Manager > Jobs.
The Job Summary page is displayed, and any active jobs are identified in the Status column as Running.
2. Click Details in the row that corresponds to the job that you want to stop.
The Job Details page is displayed.
3. Click Stop.
The selected job is stopped.
Note: Not all job types can be stopped.

Archiving Monitoring Jobs

You can archive any monitoring job with a status of Completed.

To archive a completed monitoring job:

1. Click File System Manager > Jobs.
The Job Summary page is displayed.
2. Select the check box for each job that you want to archive, and click Archive.
The Jobs Summary page is redisplayed without the selected jobs, which are moved to the Archived Jobs page.

Deleting Monitoring Jobs

You can delete any monitoring job with a status of Completed.

To delete a monitoring job:

Note: Do not attempt to delete a job whose status is Running.

1. Click File System Manager > Jobs.
The Job Summary page is displayed.
2. Select the check box for each job that you want to delete, and click Delete.
The Job Summary page is redisplayed without the selected jobs.

Deleting Archived Monitoring Jobs

To delete an archived monitoring job:

1. Click File System Manager > Jobs > Archived Jobs.
The Archived Jobs page is displayed.
2. Select the check box for each job that you want to delete and click Delete.
The Archived Jobs page is redisplayed without the selected jobs.

Servicing and Troubleshooting the System

This chapter describes servicing and troubleshooting the system. It contains the following sections:

- [“Keeping the System Up to Date” on page 238](#)
- [“Servicing Hardware Components” on page 239](#)
- [“About Field-Replaceable Units” on page 240](#)
- [“Troubleshooting LEDs” on page 241](#)
- [“Troubleshooting With the Monitoring Software” on page 256](#)
- [“Responding to System Error Messages” on page 260](#)
- [“Viewing Log Events” on page 271](#)

Keeping the System Up to Date

Use the Revision Analysis features of the diagnostic and monitoring software to keep the system patches, software, and firmware up to date. Use the Revision Analysis upgrade to do the following:

- Make sure all patches, firmware, and software are up to date at system startup.
- Establish a site process for making sure patches, firmware, and software are kept current.
- Monitor the levels manually at any time.
- Set up automated upgrades.

The diagnostic and monitoring software can be enabled to configure the Sun StorEdge Remote Response service (SSRR) as a remote notification recipient. The Sun StorEdge Remote Response service is a Sun service responds to remote notification of system alarms. The service can also check patch, firmware, and software levels and notify the local administrator when an update is available. Sun Services can make updates available from the web site or on CD-ROM.

The upgradable components of the system are:

- Storage Service Processor image
- Data Services Platform firmware
- Storage Automated Diagnostic Environment software
- Controller and loop card image
- Disk image

Because components in the system depend on each other, upgrade all components to their current version whenever a patch or new version becomes available. During the upgrade procedure, check that each component has upgraded successfully. If a component's upgrade procedure fails, repeat the procedure. If this second attempt also fails, you must back out the changes on the components that were successful so that the system operates properly. Then contact Sun Services with a description of the component's failed upgrade process.

The Revision Analysis feature also allows you to back out upgrades and restore the the following system components to their previous levels:

- Operating system
- Patches
- Firmware

Servicing Hardware Components

Service Advisor provides information and procedures for replacing Sun StorEdge 6920 system hardware.

To display Service Advisor information:

1. Go to the Java Web Console and click Sun Storage Automated Diagnostic Environment.
2. In the upper right corner of the display, click Service Advisor.
The Service Advisor page is displayed.

Click the name of the procedure for which you want instructions.

About Field-Replaceable Units

For the Sun StorEdge 6920 system, field-replaceable units (FRUs) can be replaced by Sun field engineers or by Sun-trained customer administrators. The following hardware components can be replaced at the customer site:

- Arrays: The power and cooling unit and its battery, the controller card, the loop card, and the disk drive
- Data Services Platform (DSP): switch fabric card (SFC), storage resource card (SRC), management interface card (MIC), storage I/O card (SIO), management interface Personal Computer Memory Card International Association (PCMCIA) card, power supply, fan assembly, gigabit interface converter (GBIC)
- Cabinet: power cables, power sequencers, Ethernet hubs, doors, and side panels

You can also perform the following x-option procedures:

- Add and remove array capacity and disk drives
- Add and remove array controllers
- Add and remove storage arrays
- Add cabinets
- Add, remove and disable lights-out management
- Add and remove an external modem
- Add and remove storage resource card (SRCs) sets

To replace the following components, call Sun Services:

To obtain FRU replacement instructions:

1. Click Sun Storage Automated Diagnostic Environment > Service Advisor.
The system displays the Service Advisor page.
2. Select the FRU or other option for which you want to see the procedures.

Troubleshooting LEDs

This section describes information about hardware light emitting diodes (LEDs). It contains the following subsections:




- [“Interpreting LED Status Indicators” on page 241](#)
- [“Interpreting Data Services Platform \(DSP\) LEDs” on page 242](#)
- [“Interpreting Card Status LEDs” on page 245](#)

Interpreting LED Status Indicators

The Data Services Platform (DSP) has a set of light-emitting diodes (LEDs) to identify the state of the system. In addition, the management interface card (MIC), switch fabric card (SFC), storage resource card (SRC), and storage I/O (SIO) components each have their own set of LEDs.

The following table shows how the states of LEDs are represented in this section.


TABLE 6-1 States of LEDs

LED Status	Indicates
	LED is not lit (off).
	LED is lit (on).
	LED is blinking (on and off).

Interpreting Data Services Platform (DSP) LEDs

The DSP has three sets of light emitting diodes (LEDs) to show the status of the chassis, the fans, and the power supplies.

The chassis and fan LEDs are on the front of the DSP (see the following figure).









Front view of DSP

FIGURE 6-1 Chassis and Fan LEDs on the Front of the DSP

Chassis Status

The following table shows the status of the chassis LEDs.







TABLE 6-2 Status of the Chassis LEDs

Status	indicates	Action
Green Power  Yellow Attention 	No power is applied.	Check that the power cables are connected securely.
Green Power  Yellow Attention 	Power is applied and the chassis is operating normally.	No action is required.
Green Power  Yellow Attention 	The chassis or one of the cards has a failure.	The chassis requires attention.

Fan Status

The following table shows the status of the fan LEDs.

TABLE 6-3 Status of the Fan LEDs

Status	Indicates	Action
Green Power 	Yellow Attention 	No power is applied. <ul style="list-style-type: none"> • Check that the power cables are connected securely. • Check that the fan assembly is seated properly. • Check whether one of the fans or the assembly is defective.
Green Power 	Yellow Attention 	Power is applied and the chassis is operating normally. <p>No action is required.</p>
Green Power 	Yellow Attention 	One or both fans has a failure. <p>Check that each fan can operate.</p>

Power Supply Status

The power supply LEDs are on the back of the DSP (see the following figure).






FIGURE 6-2 Power Supply LEDs on the Back of the DSP

The following table shows the status of the power supply LEDs.

TABLE 6-4 Status of the Power Supply LEDs

Status	Indicates	Action
Green DC OK 	No power is applied, or the module has failed.	<ul style="list-style-type: none"> • Check that the power cables are connected securely. • Check that the power module is functioning. • Verify that the power supply is locked into its slot.
Green DC OK 	Power is applied and DC output is normal.	No action is required.

Interpreting Card Status LEDs

Light emitting diodes (LEDs) indicate the status of a number of cards, as described in the following sections:

- “Management Interface Card (MIC) Status” on page 245
- “Switch Fabric Card (SFC) Status” on page 249
- “Storage Resource Card (SRC) Status” on page 251
- “Storage Input/Output (SIO) Status” on page 252
- “Storage Input/Output (SIO) Combo Card Status” on page 254

Management Interface Card (MIC) Status

Each system has two MICs, each with nine LEDs: the Ready, Power, Redund, and Attention LEDs are on the left side of the card, and the Act1, Act 2, Link, and Remove LEDs are on the right side (see the following figure).




FIGURE 6-3 Management Interface Card (MIC) Status

The following table shows the status of the Ready, Power, Redund, and Attention LEDs on the MIC.

TABLE 6-5 Status of Ready, Power, Redund, and Attention LEDs on the MIC



































Status	Indicates	Action
Green Ready  Yellow Redund  Green Power  Yellow Attention 	The card is operating normally.	No action is required.
Green Ready  Yellow Redund  Green Power  Yellow Attention 	No power is applied.	Check that: <ul style="list-style-type: none"> • The power cables are connected securely. • The MIC is seated properly. • The MIC is not defective.
Green Ready  Yellow Redund  Green Power  Yellow Attention 	A hardware failure has occurred.	Check that the MIC: <ul style="list-style-type: none"> • Is seated properly. • Is in the proper slot: 5 or 6. • Is not defective.
Green Ready  Yellow Redund  Green Power  Yellow Attention 	A hardware failure or software failure has occurred.	Check that the MIC is: <ul style="list-style-type: none"> • Seated properly. • Not defective.
Green Ready  Yellow Redund  Green Power  Yellow Attention 	The card is booting.	No action is required.

TABLE 6-5 Status of Ready, Power, Redund, and Attention LEDs on the MIC (*Continued*)

Status	Indicates	Action
Green Ready  Yellow Redund 	Green Power  Yellow Attention 	A hardware or software failure has occurred. Check that the MIC is: <ul style="list-style-type: none"> • Seated properly. • In the proper slot: Slot 1. • The PC card is installed. • The PC card has valid code.
Green Ready  Yellow Redund 	Green Power  Yellow Attention 	The card is booting. No action is required.
Green Ready  Yellow Redund 	Green Power  Yellow Attention 	The card is operating normally, but redundancy is active. The card is operating as the alternate master. No action is required.



The following table shows the status of the Link LEDs on the MIC.

TABLE 6-6 Status of the Link LEDs on the MIC

Status	Indicates	Action
Green Link 	The card is not connected.	Check that: <ul style="list-style-type: none"> • All connections are undamaged and secure. • The cable type is correct (straight-through). The devices at both ends of the link are powered on.
Green Link 	The card is connected.	No action is required.



The following table shows the status of the Act 1 LEDs on the MIC.

TABLE 6-7 Status of the Act 1 LEDs on the MIC

Status	Indicates	Action
Green Act 1 	There is no activity on the link.	No action is required.
Green Act 1 	There is activity on the link.	No action is required.



The following table shows the status of the Act 2 LEDs on the MIC.

TABLE 6-8 Status of the Act 2 LEDs on the MIC

Status	Indicates	Action
Green Act 2 	No Personal Computer Memory Card International Association (PCMCIA) card is installed.	No action is required.
Green Act 2 	A PCMCIA card is inserted in the slot.	No action is required.

The following table shows the status of the Remove LEDs on the MIC.

TABLE 6-9 Status of the Remove LEDs on the MIC

Status	Indicates	Action
Blue Remove 	The card is active.	Do not remove the card.
Blue Remove 	The card is not active.	The card can be removed.

Switch Fabric Card (SFC) Status

The SFC has four LEDs: Ready, Power, Redundancy, and Attention (see the following figure).




FIGURE 6-4 SFC LEDs

The following table shows the status of the LEDs on the SFC.

TABLE 6-10 Status of the LEDs on the SFC

































Status	Indicates	Action
Green Power  Yellow Attention 	Green Ready  Yellow Redundancy 	The card is operating normally. No action is required.
Green Power  Yellow Attention 	Green Ready  Yellow Redundancy 	No power is applied. Check that: <ul style="list-style-type: none"> • The power cables are connected securely. • The SFC is seated properly. • The SFC is not defective
Green Power  Yellow Attention 	Green Ready  Yellow Redundancy 	A hardware failure has occurred. Check that the card is: <ul style="list-style-type: none"> • Seated properly. • Not defective.

TABLE 6-10 Status of the LEDs on the SFC (Continued)

Status	Indicates	Action
Green Power  Yellow Attention  Green Ready  Yellow Redundancy 	<p>A hardware or software failure has occurred</p>	<ol style="list-style-type: none"> 1. Check the Event Log. See “Displaying the Event Log” on page 212. 2. Check that the card is: <ul style="list-style-type: none"> • Seated properly. • Not defective.
Green Power  Yellow Attention  Green Ready  Yellow Redundancy 	<p>The card is booting.</p>	<p>No action is required.</p>
Green Power  Yellow Attention  Green Ready  Yellow Redundancy 	<p>A hardware or software failure has occurred.</p>	<ul style="list-style-type: none"> • Check the Event Log. • See “Displaying the Event Log” on page 212.
Green Power  Yellow Attention  Green Ready  Yellow Redundancy 	<p>The card is booting.</p>	<p>No action is required.</p>
Green Ready  Yellow Attention  Green Power  Yellow Redundancy 	<p>Normal operation, but redundancy is active. The card is operating as the alternate master.</p>	<p>No action is required.</p>

Storage Resource Card (SRC) Status

The SRC has four LEDs: Ready, Power, Sync, and Attention (see the following figure).




FIGURE 6-5 SRC LEDs

The following table shows the status of the LEDs on the SRC.

TABLE 6-11 Status of the LEDs on the SRC

























Status	Indicates	Action
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	The card is operating normally. No action is required.
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	No power is applied. Check that the card is: <ul style="list-style-type: none"> Seated properly. Not defective.
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	A hardware failure has occurred. Check that the card is: <ul style="list-style-type: none"> Seated properly. Not defective.

TABLE 6-11 Status of the LEDs on the SRC (*Continued*)

Status	Indicates	Action
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	A hardware or software failure has occurred. 1. Check the Event Log. See "Displaying the Event Log" on page 212. 2. Check that the card is: <ul style="list-style-type: none"> • Seated properly. • Not defective.
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	A hardware or software failure has occurred. <ul style="list-style-type: none"> • Check the Event Log. See "Displaying the Event Log" on page 212.
Green Power  Yellow Attention 	Green Ready  Yellow Sync 	The card is booting. No action is required.

Storage Input/Output (SIO) Status

The SIO card contains two LEDs for each Gigabit interface converter (GBIC): Link and Act (see the following figure).






FIGURE 6-6 The Link and Act LEDs



The following table shows the status of the Link LEDs on the SIO card.

TABLE 6-12 Status of the Link LEDs on the SIO Card

Status	Indicates	Action
Green Link 	The card is not connected.	Check that: <ul style="list-style-type: none">• All connections are undamaged and secure.• Cable type is correct (straight-through).• Devices at both ends of the link are powered on.
Green Link 	The card is connected.	No action is required.

The following table shows the status of the Act LEDs on the SIO card.

TABLE 6-13 Status of the Act LEDs on the SIO Card

Status	Indicates	Action
Yellow Act 	There is no activity on the link.	No action is required.
Yellow Act 	There is activity on the link.	No action is required.

Storage Input/Output (SIO) Combo Card Status

The SIO Combo card contains two LEDs for each GBIC: Link and Act (see the following figure).






FIGURE 6-7 Link and Act LEDs on the SIO Combo Card



The following table shows the status of the Link LEDs on the SIO Combo card.

TABLE 6-14 Status of the Link LEDs on the SIO Combo Card

Status	Indicates	Action
Green Link 	The card is not connected.	Check that: <ul style="list-style-type: none"> • All connections are undamaged and secure. • The cable type is correct (straight-through). • The devices at both ends of the link are powered on.
Green Link 	The card is connected.	No action is required.

The following table shows the status of the Act LEDs on the SIO Combo card..

TABLE 6-15 Status of the Act LEDs on the SIO Combo Card

Status	Indicates	Action
Yellow Act 	There is no activity on the link.	No action is required.
Yellow Act 	There is no activity on the link.	No action is required.

Troubleshooting With the Monitoring Software

This section describes troubleshooting procedures. It contains the following subsections:

- [“Testing an Email Address” on page 256](#)
- [“Generating a Solution Extract” on page 257](#)
- [“Displaying a Solution Extract” on page 257](#)
- [“Working With Core Files” on page 258](#)
- [“Running a Diagnostic Test on a Device” on page 259](#)

Testing an Email Address

To verify that the monitoring and diagnostic software can successfully send email to a specific email address:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Notification.

The Notification Setup page is displayed.

2. Click Test Email.

The Send Test Email window is displayed.

3. In the To field, specify the email address you want to test.
4. (Optional) Specify a test message in the Message field.
5. Click Send.

If the message is successfully sent, the following message is displayed:

```
The test email has been sent successfully.
```

6. Click Close to close the Send Test Email window.

Generating a Solution Extract

You can use the Solution Extract function to generate a snapshot of the Sun StorEdge 6920 system configuration.

To generate a solution extract snapshot:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Utilities > Solution Extract.

The Solution Extract page is displayed.

2. Click Extract.

A .tar file is created containing the snapshot of the Sun Solution configuration.

Note: The extract report may take several minutes to run. You can monitor the progress of the solution extract using the Jobs function.

3. Do any of the following:

- To display the log file for the extraction process, click View Log.

The extraction log file is displayed.

- To download the extracted .tar file, click Download.

- To display the contents of the Solution Extract, click View Content. The View Content option appears only if you are logged in as root.

Note: The contents of a solution extract are not formatted for displaying without specific knowledge of extract data formats.

Displaying a Solution Extract

After generating a solution extract of the Sun Solution configuration, you can use the Solution Extract viewer to display the extract data.

Note: The View Content option appears only if you are logged in as root. The contents of a solution extract are not formatted for displaying without specific knowledge of extract data.

To display a solution extract snapshot:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Utilities > Solution Extract.

The Solution Extract page is displayed.

2. In the Actions column, click View Content for the solution extract you want to display.

Note: If there are no solution extracts listed, you have to first generate a solution extract.

Working With Core Files

This section describes core files. It contains the following subsections:

- [“About Core Files” on page 258](#)
- [“Downloading and Deleting Core Files” on page 258](#)

About Core Files

A core file is a memory dump of the current state of a device, when the device is unable to boot or unable to continue running. In the event of a crash, the device puts its core file on the Storage Service Processor. The system software retains up to five core files for each device. Special tools can be used to analyze a core file and debug the problem.

You can download core files for analysis and troubleshooting, and you can delete selected core files.

Note: Each new core file from a device overwrites any preceding core file.

Downloading and Deleting Core Files

To download or delete a core file:

1. Click Sun Storage Automated Diagnostic Environment > Administration > Utilities > Core Files.

The Core Files page displays any existing core files by device. A maximum of five core files can be displayed for each device.

2. Do one of the following:
 - To download a core file, click the name of the core file you want to download.
Note: If your browser does not open a dialog box to initiate the download, press the Shift key while clicking the file name to download the file.

- To delete a core file, click the check box in the Delete column that corresponds to the core file you want to delete, and click Delete Selected.

Running a Diagnostic Test on a Device

To run a diagnostic test on a device:

1. Click Sun Storage Automated Diagnostic Environment > Inventory.
The Devices page displays summary information for each system device.
2. Click the component name of the device for which you want to run a diagnostic test.
Details for the selected device are displayed.
3. Click Diagnostics.
The Device Details page scrolls to a list of available diagnostic tests for the device.
4. Select the desired diagnostic test, and click Run Test.
The Diagnostic Setup page is displayed for the selected diagnostic test.
5. Complete the diagnostic setup options and click Run.
The Test Results page is displayed. Results are added when the test is completed.

Responding to System Error Messages

This topic lists some common system error messages and recommends the action to take for each message.

A missing volume cannot be modified.

Explanation: If the virtual disks that make up a volume become degraded, the software lists the volume's status as Missing. You can delete, but not modify, the volume.

Action: Do the following:

- Resolve all Down and Critical system alarms that are displayed in the page banner.
- Make sure the status of the virtual disks is Online.
- Modify the volume when the software displays its status as OK.

A sync must be performed before the roles can be reversed.

Explanation: Before you can perform a role reversal, the primary and secondary volumes within a replication set must be synchronized.

Action: Resume replication to synchronize the volumes. When the software displays the replication progress status as Replicating, perform the role reversal procedure.

At least two virtual disks must exist in the system to create a legacy volume.

Not enough initialized Virtual Disks for legacy volume metadata.

Explanation: When you create a legacy volume from an external storage device, the system requires at least two virtual disks to have been created in any storage pool. The metadata that manages all legacy volumes is not kept on the legacy volume but is kept redundantly in a reserved area on other virtual disks.

Action: Do the following:

1. From any system pool's Storage Pool Details page, click the Add Storage to Pool button to do one of the following:
 - Create new virtual disks from storage arrays
 - Move unused virtual disks from other pools into this pool
2. Try to create or modify the legacy volume again.

You can see how many virtual disks exist in the system on the Virtual Disk Summary page.

Description contains too many characters.

Explanation: You exceeded the limit of 256 characters.

Action: Enter a description with fewer than 256 characters.

Description is a required field.

Explanation: The field is empty.

Action: Enter identifying text in the Description field.

Mirror creation failed. The following errors occurred: Volume creation failed. The following errors occurred: The existing overlap failed. Array isolation is impossible.

Explanation: A volume or mirror could not be created. Additional error information provides details that explain why the operation failed. In some cases, the additional error information includes multiple reasons for the failure.

Action: Do the following:

- Correct one of the errors and try to create the volume or mirror again.

- If you want to create, or add to, a mirrored volume with an isolation policy of Required Isolation, you might need to add storage to one or more of the pools from which you are creating the mirror; the storage you add must come from a storage array different from those already in the pools.
- From the Pool Summary page, use the Storage Pool by Array Assignment view to determine which pools have virtual disks from more than one array. You can use these pools to create or add to a mirror with the Required Isolation isolation policy.
- If you do not require the mirror components to be isolated on separate arrays, which reduces their tolerance to a failure, select the Isolation Optional isolation policy.

Name contains too many characters.

Explanation: You exceeded the limit of 8 characters.

Action: Enter a name using 8 characters or fewer.

Name is a required field.

Explanation: The field is empty.

Action: Enter identifying text in the Name field.

Name server is not a valid IP address.

Explanation: You entered an invalid value.

Action: Check the IP address of your network's domain name server.

No items were selected on which to perform the action.

Explanation: You did not select an element on which to perform the action.

Action: Select an element first, and then select the action.

NTP Server IP address must be specified when using an NTP server.

Explanation: You did not enable Network Time Protocol (NTP) correctly.

Action: To enable NTP, clear the check box and type the IP address of your network's NTP server.

Number of volumes exceeds maximum allowed.

Explanation: You exceeded the number of volumes that can be created.

Action: Delete one or more existing volumes. Create the new volumes again.

Requested size exceeds available capacity.

Explanation: You exceeded the available capacity.

Action: Do the following:

1. From any system pool's Storage Pool Details page, click the Add Storage to Pool button to do one of the following:
 - Create new virtual disks from storage arrays
 - Move unused virtual disks from other pools into this pool
2. Try to perform the operation again.

You can see how many virtual disks exist in the system on the Virtual Disk Summary page.

Requested size must be between 1 and *variable*.

Explanation: You entered a value outside the defined range.

Action: Enter a value this is within the range displayed in the error message.

Specified gateway IP address is not a valid IP address.

Explanation: You entered an invalid value.

Action: Check the IP address for the system that serves as your network's gateway.

Specified IP address is not valid.

Explanation: You entered an invalid value.

Action: Check the format of the IP address for a missing number or period.

Specified netmask is not a valid IP address.

Explanation: You entered an invalid value.

Action: Check the format of the netmask.

Synchronization cannot be started because the remote volume is not secondary.

Explanation: The replication set on the remote peer is configured incorrectly. For successful replication, the replication set must be configured with a primary volume on one peer and a secondary volume on the other peer.

Action: Change the role of the replication set (or of the consistency group of which it is a member) to secondary on the remote peer and resume replication on the primary replication set.

The Array type specified is not valid.

Explanation: If a profile applies a value that is inappropriate for the array type, you cannot add storage to the pool.

Action: Check the array type defined for the storage pool.

The number of Disks specified is not valid.

Explanation: You entered an invalid value.

Action: If a pool has a number defined, the number of disks in a virtual disk must be the same for all virtual disks in the pool.

The operation could not take place because the bitmap is being set or cleared.

Explanation: This is a temporary condition. The system is operating on the replication bitmap and cannot perform the operation.

Action: Wait briefly and retry the operation.

The PIT for sync startup failed.

Explanation: The secondary peer could not take a snapshot and the resume operation failed.

When replication resumes, if the replication set's role is secondary and the Snap before Synchronization option is enabled, the secondary peer automatically takes a snapshot before synchronization starts and deletes the snapshot when synchronization is complete.

Action: Make sure that the secondary volume has been configured with sufficient snapshot reserve space and is under the limit of active snapshots (8). Then resume replication on the primary replication set or consistency group.

The RAID level specified is not valid.

Explanation: The RAID levels among storage trays must match.

Action: Make sure the RAID levels match.

The resulting mirrored volume would leave no readable components.

Explanation: The mirrored volume operation might result in a mirror with no component that is capable of providing the volume's data. At least one mirror component must be in a valid state to satisfy I/O requests for the volume's data.

Action: If split components that are currently resilvering cannot provide the volume's data, wait for the resilvering operation to complete and try the operation again.

The segment size specified is not valid.

Explanation: The segment sizes must match.

Action: Specify a different segment size.

The stripe size specified is not valid.

Explanation: Valid sizes are 128 KB, 512 KB, 1 MB, 2 MB, 4 MB, 8 MB, and 16 MB.

Action: Specify a valid stripe size.

The virtualization strategy specified is not valid.

Explanation: The volumes in a storage pool must use the same virtualization strategy: striping or concatenation.

Action: Check the minimum number of virtual disks for that strategy.

There are not enough arrays to isolate logs.

Explanation: The system cannot find two virtual disks from separate storage arrays with sufficient free space to provision the mirror logs.

The mirror logs, which are used to track changes to the mirror so components remain synchronized, must always be located on two different virtual disks. If the isolation policy is Required Isolation, the two virtual disks must be on separate, or isolated, storage arrays.

The virtual disks selected for mirror logs are not required to be on storage arrays isolated from the mirror components; however, the system chooses virtual disks isolated from the mirror components whenever possible.

If you want to create a two-component mirrored volume with an isolation policy of Required Isolation, the pools you choose must contain at least two virtual disks from two different storage arrays. Similarly, a three-component mirrored volume requires at least three virtual disks from three different storage arrays.

Action: Do the following:

- Add additional storage to the pools from which you created the mirror. From the pool's Details page, click the Add Storage to Pool button.
- From the Pool Summary page, use the Storage Pool by Array Assignment view to determine which pools have virtual disks from more than one array and can be used to create or add to a mirror whose isolation policy is Required Isolation.

If you do not require the mirror components to be isolated on separate arrays (reducing their tolerance to a failure), you can use the Optional Isolation policy. When you create a mirrored volume from an existing volume or add a component from an existing volume, the system includes the storage pool of the existing volume when provisioning these mirror logs.

Time should be in the format: hh:mm.

Explanation: The time format is incorrect.

Action: Enter the time again using the *hh:mm* format.

Unexpected internal system error. Retry the operation and consult the Service Advisor in Storage Automated Diagnostic Environment if the error persists.

Explanation: An internal system error has occurred.

Action: Resolve all Down and Critical system alarms that are displayed in the page banner and try the operation again. For information, go to the Sun Java™ Web Console page and click Sun Storage Automated Diagnostic Environment > Service Advisor > X-Options.

User authentication failed. Please try again.

Explanation: The password is invalid. If you have lost or forgotten the password, contact Sun Services.

Action: Enter the correct password for the account.

User *variable* is not allowed to access Management software; please log out and log in as a valid user.

Explanation: You are logged in using the incorrect role.

Action: Log out and log in again using the storage or admin role.

Volume size too small (should be over 10MB).

Explanation: You entered an invalid value.

Action: Enter a value greater than 10 MB.

Volumes whose virtualization type is stripe require a minimum of two virtual disks in storage pools.

Explanation: When you tried to create or modify a volume or mirror, one or more of the pools that you specified did not have sufficient virtual disks with available capacity to complete the configuration.

Action: Do the following:

- Navigate to the Storage Pool Details page of the pools you specified and click the Add Storage to Pool button. Either create new virtual disks from storage arrays or move unused virtual disks from other pools into this pool. Then try to create or modify the volume or mirror again.
- Use the Virtual Disk Summary page to view virtual disks and their available capacity, and the storage pools in which they reside. Virtual disks have a fixed number of partitions for volume data (see the Virtual Disk Details page). If a virtual disk has available capacity but user volumes consume all the partitions, the available capacity cannot be used for volume data. The current in-use partition count for each virtual disk can also be seen on the Virtual Disk Summary page.

Note: Storage for a mirrored volume's mirror logs comes from any pool that you select to create the mirrored volume; you do not specify a separate storage pool for the mirror logs. Therefore, the pools you select must have sufficient virtual disks with sufficient available capacity to include the mirror logs. Each mirror log is 2 Mbytes in size; each mirrored volume requires two mirror logs that must be provisioned on separate virtual disks, as well as one mirror log for each component. That is, the combination of pools used to create a two component mirrored volume requires at least two virtual disks with available capacity greater than 2 Mbytes, and two (the same or different) virtual disks with available capacity for two more mirror logs of 2 Mbytes each. When you create a mirrored volume from an existing volume or add a component from an existing volume, the system includes the storage pool of the existing volume when provisioning these mirror logs.

This explanation and action also apply to the following error messages:

- **A volume whose virtualization type is stripe requires that you specify a storage pool with at least two virtual disks.**
- **BV was unable to provision the mirror components.**
- **Not enough virtual disks are available for log data.**
- **One or more of the specified virtual devices are too small.**
- **requested size exceeds available capacity.**
- **The maximum number of partitions on this virtual disk has been reached.**
- **The maximum number of partitions per virtual disk has been reached.**
- **The mirror log's storage pool has too few virtual disks.**

- The volume configuration operation failed: All virtual disks exceed the partition limit.
- The volume size specified is too large for the virtual disks in the storage pool specified.
- The virtual disk(s) in the storage pool specified have insufficient free space.
- The virtual disk size or sizes in the storage pool specified are not large enough for the volume.
- Volumes whose virtualization type is stripe require a minimum of two virtual disks with available free space.

Viewing Log Events

The system logs its events to the same log file maintained by the Java Web Console.

To view log events:

1. Click Sun StorEdge 6920 Configuration Service > Administration > Notification Management.
2. Click the indicated link to view log events.

Alternatively, click Sun Storage Automated Diagnostic Environment > Administration > Event Log.

Glossary

Definitions obtained from the Storage Networking Industry Association (SNIA) Dictionary are indicated with “(SNIA)” at the end. For the complete SNIA Dictionary, go to www.snia.org/education/dictionary.

- agent** The component of the system monitoring and diagnostic software that collects health and asset information about the system.
- alarm** A type of event that requires service action. See also event.
- array** Multiple disk drives that function as a single storage device. A high-availability (HA) array configuration has multiple controller and expansion trays of disk drives. See also dedicated hot-spare and hot-spare.
- array hot-spare** A spare disk within an array that is used for failover when any of the active disks fail. See also dedicated hot-spare and hot-spare.
- array type** An internal storage array configuration that is defined by the number of controller units and the total number of storage trays. For example, a 2x4 storage array configuration consists of a total of four storage trays, two of which are controller trays.
- asynchronous queue** In the context of data replication, a queue used to store writes that are to be replicated to the remote site. After the writes have been put into the queue, the writes are acknowledged to the application and then forwarded to the remote site as network capabilities permit. The asynchronous queue is a persistent queue, so in the event of a disaster at the primary site, the data in the asynchronous queue is not lost.
- asynchronous replication** A form of data replication in which application write operations are written to the primary site and to the asynchronous queue on the primary site. The asynchronous queue forwards queued writes to the secondary site as network capabilities permit. The write operations to the primary site are confirmed,

regardless of when, or whether, they are replicated successfully to the secondary site. Deferring the secondary copy removes long-distance propagation delays from the I/O response time. See also synchronous replication.

- autosynchronization** An option enabled at the primary site that attempts to synchronize replication sets or consistency groups whenever a link is established. With autosynchronization, synchronization continues even if there are link errors, for example.
- base cabinet** The system's main cabinet, which contains a Data Services Platform (DSP), storage trays, a Storage Service Processor with a USB flash disk, a Storage Service Processor accessory tray, a service panel, and power distribution units (PDUs). The base cabinet is pre-wired with Ethernet, Fibre Channel, and power cables and can be connected to one or two expansion cabinets. See also expansion cabinet.
- break** To remove a mirror component from the mirror, voiding its relationship with the other mirror components. The broken mirror component becomes a standalone volume in the system, and synchronization with the other mirror components is not maintained. See also mirror component and split.
- captive storage** See internal storage.
- combo card** See storage I/O (SIO) card.
- component** See mirror component.
- concatenation** A storage allocation method in which sequential blocks on a disk are linked together to form a larger logical device. This method combines the storage potential of several physical devices. See also striping.
- consistency group** A collection of replication sets grouped together to ensure write order consistency across all the replication sets' primary volumes. An operation on a consistency group applies to all the replication sets within the consistency group, and consequently their volumes.
- controller pair** A pair of controller units that services a group of storage trays.
- controller tray** A storage tray with an installed RAID controller and up to 14 disk drives. In a Sun StorEdge 6920 system, a pair of controller trays is the smallest possible storage array configuration, the 2x2 array type. See also expansion tray.
- controller unit** The card that manages RAID functions and failover characteristics for an array.
- control path** The route used for communication of system management information, usually an out-of-band connection. See also out-of-band management.

copy on write	A technique for maintaining a point in time copy of a collection of data by copying only data which is modified after the instant of replicate initiation. The original source data is used to satisfy read requests for both the source data itself and for the unmodified portion of the point in time copy. (SNIA) See also snapshot reserve space.
customer LAN	See site LAN.
DAS	See direct attached storage (DAS).
data host	Any host that uses the system for storage. A data host can be connected directly to the system (direct attached storage, or DAS) or can be connected to an external switch that supports multiple data hosts (storage area network, or SAN).
data path	The route taken by a data packet between a data host and the storage array. See also in-band management.
data replication	A disaster recovery and business continuance method in which a primary volume at the local site and a secondary volume at a remote site contain the same data on an ongoing basis, thereby protecting user data.
Data Services Platform (DSP)	The controller component of the Sun StorEdge 6920 system, which consolidates and virtualizes storage so that all storage in the system can be managed as a single scalable entity.
data striping	See striping.
dedicated hot-spare	A disk that serves as the hot-spare for one and only one virtual disk in a storage tray.
degraded	The condition of a volume in which one or more input or output data paths are not operating properly. Although the redundant failover paths are still intact, a degraded volume holds no significant value for the storage configuration and should therefore probably be deleted from the system.
dependent copy	See snapshot.
direct attached storage (DAS)	A storage architecture in which the systems that store data are connected physically to storage elements.
disk	A physical, non-volatile, rewritable data storage device. See also virtual disk.
domain	See storage domain.
DSP	See Data Services Platform (DSP).
electro-magnetic interference	Radiated electro-magnetic signals that can interfere with the normal transmission of information.

- EMI** See electro-magnetic interference.
- event** Any condition reported by a device to the system monitoring agent. See also alarm.
- expansion cabinet** A cabinet, pre-wired with Ethernet, Fibre Channel, and power cables, connected to the base cabinet to increase storage capacity. A power management cable connects the service panels of the cabinets, and the Fibre Channel/Ethernet cable connects the I/O panels of the cabinets. See also base cabinet.
- expansion tray** A storage tray that has up to 14 disk drives, but does not have a RAID controller. This type of storage tray is used to expand the capacity of an array and must be attached to a controller tray to function. See also controller tray.
- extend** In the context of managing volumes, to increase a volume's capacity.
- extent** A set of contiguous blocks with consecutive logical addresses on a physical or virtual disk.
- external storage** A physical disk or storage array connected to the Sun StorEdge 6920 system located outside of the base or expansion cabinets.
- fabric** A Fibre Channel switch or two or more Fibre Channel switches interconnected in such a way that data can be physically transmitted between any two N_Ports on any of the switches. (SNIA)
- failover and recovery** The process of changing the data path automatically to an alternate path.
- Fast Start operation** An option of the suspend operation and the procedure in which a method such as a backup tape is used to copy data from the primary volume to the secondary volume. This procedure is used to avoid the initial step in which you send the primary volume's data over the physical link. For example, network bandwidth might justify a Fast Start procedure. See also resume operation and suspend operation.
- fault signature analysis** An algorithm applied by the diagnostic and monitoring software to a specific set of events that can be correlated by time and locality to a specific cause. Fault signature analysis assumes that the most significant event is the most probable cause and aggregates the remaining events underneath that event to improve the signal/noise ratio when a single probable cause might result in a multitude of events.
- FC** See Fibre Channel (FC).
- FC port** See Fibre Channel (FC) port.
- FC switch** See Fibre Channel (FC) switch.
- Fibre Channel (FC)** A set of standards for a serial I/O bus capable of transferring data between two ports at up to 100 MBytes/second, with standards proposals to go to higher speeds. Fibre Channel supports point to point, arbitrated loop, and

switched topologies. Fibre Channel was completely developed through industry cooperation, unlike SCSI, which was developed by a vendor and submitted for standardization after the fact. (SNIA)

- Fibre Channel (FC) port** A port on the I/O panel that connects data hosts, external storage, or internal storage to the Sun StorEdge 6920 system. See also host port and storage port.
- Fibre Channel (FC) switch** A networking device that can send packets directly to a port associated with a given network address in a Fibre Channel storage area network (SAN). Fibre Channel switches can be used to expand the number of data host or external storage device connections. Each switch is managed by its own management software.
- field-replaceable unit (FRU)** An assembly component that is designed to be replaced on site, without the system having to be returned to the manufacturer for repair.
- flash card or disk** See USB flash disk.
- force break** To remove a mirror component (mirrored volume) from the mirror before the resilvering process is complete, causing the condition of its data to be degraded. Just as with a break operation, the mirror component for which you implement a force break becomes a standalone volume in the system; however, it is inaccessible and should probably be deleted. See also break, degraded, mirror component, and resilvering.
- free** The state of a volume that is not mapped to initiators.
- FRU** See field-replaceable unit (FRU).
- FSA** See fault signature analysis.
- full synchronization** A resume operation in which a complete volume-to-volume copy occurs. Unlike a normal resume operation, in which a copy of differences between the primary and secondary volumes occurs, a full synchronize operation copies the entire contents of the volume. The system performs a full synchronize operation the first time you resume data replication on a replication set. See resume operation and synchronization.
- HBA** See host bus adapter (HBA).
- host bus adapter (HBA)** An I/O adapter than connects a host I/O bus to a computer's memory system. (SNIA) See also initiator.
- host port** A port on the I/O panel that connects to a data host. See also storage port.
- hot-spare** The drive used by a controller to replace a failed disk. See also array hot-spare and dedicated hot-spare.

- in-band management** Transmission of a protocol other than the primary data protocol over the same medium as the primary data protocol. Management protocols are a common example of in-band transmission. (SNIA)
- The Sun StorEdge 6920 system uses an in-band management path between hosts and the storage arrays to transport both data and management traffic. See also out-of-band management.
- independent copy** See mirror component and mirroring.
- initiator** A system component that initiates an I/O operation over a Fibre Channel (FC) network. If allowed by FC fabric zoning rules, each host connection within the FC network has the ability to initiate transactions with the storage array. Each host in the FC network represents a separate initiator, so if a host is connected to the system through two host bus adapters (HBAs), the system identifies two different initiators (similar to multi-homed, Ethernet-based hosts). In contrast, when multipathing is used in round-robin mode, multiple HBAs are grouped together, and the multipathing software identifies the group of HBAs as a single initiator.
- internal storage** An array physically housed in the Sun StorEdge 6920 system base cabinet, or expansion cabinet, and managed by the system management software. See also external storage.
- I/O panel** The portion of the service panel that provides Fibre Channel (FC) port connections for data hosts, internal and external storage, gigabit Ethernet ports for remote replication, power connections for up to two expansion cabinets, and Ethernet and FC connections for expansion cabinets. See also service panel and Service Processor panel.
- IOPS** A measure of transaction speed, representing the number of input and output transactions per second.
- LAN** See local area network (LAN).
- legacy volume** An entire LUN on an external storage array that you can use in specific ways as if it were any other local volume, while preserving the user data on that external storage array. You can apply the system's data services to a legacy volume, however, you cannot extend a legacy volume.
- local area network (LAN)** A communications infrastructure designed to use dedicated wiring over a limited distance (typically a diameter of less than five kilometers) to connect a large number of intercommunicating nodes.
- logical unit number (LUN)** The SCSI identifier of a logical unit with a target. (SNIA)
- In the Sun StorEdge 6920 system, a LUN is the number that is associated with mapping of a volume to an initiator.
- LUN** See logical unit number (LUN).

LUN mapping	The process by which volume permissions (read only, read/write, or none) are assigned to an initiator.
LUN masking	The process by which a mapped initiator is filtered.
MAC address	See media access control (MAC) address.
management host	A host with an in-band and/or out-of-band network connection to the system that is used to manage the system. A management host can have monitoring software installed, such as a remote CLI package, Sun StorEdge Enterprise Storage Manager, Sun Storage Automated Diagnostic Environment (Enterprise Edition), or a third-party monitoring program. A management host can also be used to run management software on another machine through a network connection.
management interface card (MIC)	The card on which the management software resides. Each Data Services Platform has two MICs, which operate in a master/alternate master mode.
management path	See out-of-band management.
master agent	In Sun Storage Automated Diagnostic Environment, Enterprise Edition, the primary health and monitoring agent designated by the user to act as the aggregation point for one or more slave agents. The responsibilities of the master agent include analyzing the events forwarded from the slave agents, generating alarm notification to local and remote recipients, and provisioning the user interface to the slave agents. See also master / alternate master and slave agent.
master / alternate master	In Sun Storage Automated Diagnostic Environment, Enterprise Edition, a slave agent designated by the user to act as the temporary master agent if the master agent fails to send a heartbeat signal to the alternate master during a defined time period. The alternate master does not assume all the responsibilities of the master agent; it simply acts as the aggregation point for the remaining slave agents, collecting their events, until the master agent returns to active service. See also master agent and slave agent.
media access control (MAC) address	The physical address identifying an Ethernet controller board. The MAC address, also called an Ethernet address, is set at the factory and must be mapped to the IP address of the device.
MIC	See management interface card (MIC).
mirror	A special type of volume in the Sun StorEdge 6920 system, consisting of up to four separate yet equal mirror components you can access independently (through the mirror) and use to track changes to and update your data set and manage your data migration strategy.

mirror component One of up to four individual copies of the same data set that constitute a mirror in the Sun StorEdge 6920 system. When you perform an action on a mirror component, the resilvering process synchronizes the mirror so that each component is an equal yet separate copy of the same data set. You can perform a variety of operations on a mirror component. See also break, force break, rejoin, reverse rejoin, resilvering, and split.

mirrored volume See mirror and mirror component.

mirroring A form of storage – also called RAID Level 1, independent copy, and real-time copy – whereby two or more independent, identical copies of data are maintained on separate media. Typical mirroring technologies enable the cloning of data sets to provide redundancy for a storage system. The Sun StorEdge 6920 system enables you to create and manipulate up to four equal mirror components for each mirror and to access each component individually in order to track changes to and preserve the integrity of your data. See also mirror and mirror component.

mirror log Area of the storage pool used to track the state of mirror components (such as resilvering progress) relative to the mirror as a whole.

multipathing A design for redundancy that provides at least two physical paths to a target.

**Network Storage
Command Center
(NSCC)**

A repository and application comprised of a database and a browser-based user interface, designed and maintained by Sun, to collect and analyze health and performance data from supported storage devices.

non-captive storage See external storage.

NSCC See Network Storage Command Center (NSCC).

notification The process performed by the master agent when one or more events require fault reporting to a configured notification recipient such as a local email address, an SNMP port, or a remote service such as the Sun StorEdge Remote Response service. The master agent constructs the event and assembles the information necessary to transmit the event by way of the user-configured transport mechanism. The monitoring and diagnostic software supports notification of one or more notification recipients to satisfy the desired level of notification. Note: The Sun StorEdge 6920 array contains a single agent, which effectively acts as the master agent for the array.

original volume The starting point for a mirrored volume or data migration. In the context of mirroring, see also primary component.

**out-of-band
management**

Transmission of management information for Fibre Channel components outside of the Fibre Channel network, typically over an Ethernet network. (SNIA) The Sun StorEdge 6920 system is managed over an out-of-band network by way of an Ethernet connection between the service panel and the local area network (LAN). See also in-band management.

parallel monitoring	A monitoring control that allows the agent to examine more than a single device in parallel during an agent polling cycle. The positive effect of increasing the value for this control beyond the default (1) is to speed up the monitoring of devices when a large population exists. The negative effect of increasing the value of this control is that the agents will consume more memory and CPU cycles.
parent volume	The volume for which a snapshot is taken. See also snapshot.
patch	A software or firmware update for a storage device or device component.
PDU	See power distribution unit (PDU).
point in time copy	A fully usable copy of a defined collection of data that contains an image of the data as it appeared at a single point in time. The copy is considered to have logically occurred at that point in time, but implementations may perform part or all of the copy at other times (e.g., via database log replay or rollback) as long as the result is a consistent copy of the data as it appeared at that point in time. Implementations may restrict point in time copies to be read-only or may permit subsequent writes to the copy. (SNIA) In the context of data replication, the storage system that houses the primary replication set, which is the remote counterpart of the secondary site. Also known as the local site copy. See also snapshot.
policy	A rule or guideline that can result in an automatic response to a system event.
pool	See storage pool.
power distribution unit (PDU)	The assembly that provides power management for the system. The Sun StorEdge 6920 system provides two PDUs that enable the management software to control the distribution of power to system components for Lights-Out Management (LOM) and field-replaceable unit (FRU) service actions.
primary peer	One of a pair of physically separate systems on which the primary replication set resides. The primary peer copies user data to its counterpart, which is the remote, secondary peer.
primary component	The first component created to establish a mirrored volume. See also mirror, mirror component, and mirrored volume.
primary volume	The volume that contains the original user data that the primary peer replicates to the secondary peer.
profile	See storage profile.
provisioning	The process of allocation and assignment of storage to hosts.
RAID	An acronym for Redundant Array of Independent Disks, a family of techniques for managing multiple disks to deliver desirable cost, data availability, and performance characteristics to host environments. (SNIA)

real-time copy	See mirroring.
reconstruction	The process of rebuilding of lost data on a replacement disk after a disk failure.
rejoin	To move a split mirror component back into the mirror such that, when the resilvering process is complete, the mirror component has been made identical to all other mirror components in the mirror. See also mirror component, reverse rejoin, and split.
remote scripting CLI client	A command-line interface (CLI) that enables you to manage the system from a remote management host. The client communicates with the management software through a secure out-of-band interface, HTTPS, and provides the same control and monitoring capability as the browser interface. The client must be installed on a host that has network access to the system.
replication	See data replication.
replication bitmap	The bitmap that tracks changes to the primary volume. Writes issued to the primary peer are noted in the replication bitmap. The replication set at the secondary peer also includes a replication bitmap that tracks changes if a role reversal assigns the secondary volume the role of primary.
replication link	A logical connection associated with a Gigabit Ethernet port that transports data and replication control commands between primary and secondary sites. The Gigabit Ethernet ports at both sites must be enabled for data replication and be configured with the remote site's IP information.
replication peer	One of a pair of complimentary components that are on physically separate systems. For example, user data is copied to a remote system, which is the counterpart, or remote peer, of the system on which that user data resides.
replication set	A local volume paired with reference to a single remote volume on a remote peer. A replication set works in conjunction with an identically configured replication set on the remote peer to provide an instance of replication. The local volume within a replication set is associated with a replication bitmap and, depending on the set's attributes, with an asynchronous queue.
resilvering	The synchronization of mirror components such that, when the process is complete, the mirror comprises equal, independent copies of the same data set. Resilvering occurs when you have performed a rejoin or reverse rejoin operation on a mirror component.
resnap	To create a snapshot again and replace the old with the new.
resume operation	In the context of data replication, a synchronization operation to establish an identical copy of the primary volume's user data on the secondary volume. The data is synchronized when replication occurs. Synchronization can be initiated by either the user or the system. See also autosynchronization, suspend operation, and synchronization.

reverse rejoin	To move a split mirror component back into the mirror such that, when the resilvering process is complete, all mirror components in the mirror have been made identical to the previously split mirror component. See also break, mirror component, rejoin, resilvering, and split.
reverse synchronization	See role reversal.
revision analysis	In Sun Storage Automated Diagnostic Environment, Enterprise Edition, the process of collecting the current revision information for the software and firmware elements of the system and comparing them against a set of acceptable levels. See also revision maintenance.
revision backout	The removal of a patch update on a storage device or device component. See also revision upgrade.
revision maintenance	A system process that combines performing revision analysis on the elements of the system and locating, acquiring, and installing the necessary deliverables to bring elements up to an acceptable revision level. See also revision analysis.
revision upgrade	The installation of a patch update on a storage device or device component. See also revision backout.
role reversal	In the context of data replication, a procedure in which the secondary host is assigned the role of primary host within an established replication set, and the primary volume is updated with the contents of the secondary volume. Role reversal is a failover technique used when the primary site fails and for disaster rehearsal.
rollback	The process by which a volume's data is reset to become identical to a snapshot taken of that volume.
SAN	See storage area network (SAN).
secondary peer	One of a pair of physically separate systems on which the secondary replication set resides. The secondary peer is the recipient of user data from its counterpart, which is the primary peer.
secondary volume	The remote counterpart of the primary volume. The secondary volume is the replicated copy of the primary volume. You can map or create a volume snapshot of a secondary volume. You cannot read from or write to a secondary volume unless it is in scoreboard mode or you change its role to primary.
Service Advisor	A diagnostic tool component that provides tools and procedures for performing service on storage device.
service panel	A group of input and output connections located at the back of the base cabinet that provides the cabling interface to both control path functions and data path functions. The service panel consists of the Service Processor panel and the I/O panel. See also I/O panel and Service Processor panel.

**Service Processor
panel**

The portion of the service panel that provides a modem connection, LAN connections, serial ports, and an AUX port for connection of the Data Services Platform (DSP) management interface card (MIC).

SFC See switch fabric card (SFC).

**Simple Mail Transfer
Protocol (SMTP)**

A protocol for sending e-mail messages between servers. Most e-mail systems that send mail over the Internet use SMTP to send messages from one server to another; the messages can then be retrieved with an email client using either Post Office Protocol (POP) or Internet Message Access Protocol (IMAP). In addition, SMTP is generally used to send messages from a mail client to a mail server. This is why you need to specify both the POP or IMAP server and the SMTP server when you configure your email application. (Webopedia)

**Simple Network
Management Protocol
(SNMP)**

An IETF protocol for monitoring and managing systems and devices in a network. The data being monitored and managed is defined by a Management Information Base (MIB). The functions supported by the protocol are the request and retrieval of data, the setting or writing of data, and traps that signal the occurrence of events. (SNIA)

SIO card See storage I/O (SIO) card.

site LAN The local area network at your site. The system connects to your LAN through the USER LAN port on the service panel. When the system is connected to your LAN, the system can be managed through a browser from any host on the LAN.

slave agent In Sun Storage Automated Diagnostic Environment, Enterprise Edition, a health and monitoring agent that collects health and performance data from the devices it has been designated to monitor and forwards events to the master or alternate master agent for evaluation and notification processing. Slave agents cannot function completely without a master or alternate master agent. See also master agent and master / alternate master.

SMTP See Simple Mail Transfer Protocol (SMTP).

snapshot An instantaneous copy of volume data at a specific point in time. Snapshots are stored in snapshot reserve space on the (parent) volume for which they are taken.

**snapshot reserve
space**

Storage space, taken from a pool, where the system stores copies of the parent volume's original data before it is overwritten. See also storage pool.

split To separate a mirror component from the mirror, with the intent to rejoin it later. A split component is counted toward the limit of four mirror components per mirror and continues to be tracked by the system as part of the mirror. See also component, rejoin, reverse rejoin, and snapshot.

SRC	See storage resource card (SRC).
SSRR	See Sun StorEdge Remote Response service.
storage area network (SAN)	An architecture in which the storage elements are connected to each other and to a server that is the access point for all systems that use the SAN to store data.
storage domain	A secure container that holds a subset of the system's total storage resources. Multiple storage domains can be created to securely partition the system's total set of storage resources. This enables you to organize multiple departments or applications into a single storage management infrastructure.
storage I/O (SIO) card	A card that provides the Fibre Channel (FC) ports for the Data Services Platform (DSP). This card is always paired with a storage resource card (SRC). The Sun StorEdge 6920 system supports two types of SIO cards. The SIO-8 card has eight FC ports and the SIO COMBO card has six FC ports and one Gigabit Ethernet port. See also storage resource card (SRC) set.
storage pool	A container that groups physical disk capacity (abstracted as virtual disks in the browser interface) into a logical pool of available storage capacity. A storage pool's characteristics are defined by a storage profile. You can create multiple storage pools to segregate storage capacity for use in various types of applications (for example, high throughput and online transaction-processing applications).
storage port	A port on the I/O panel that connects to internal storage. See also host port.
storage profile	A defined set of storage performance characteristics such as RAID level, segment size, dedicated hot-spare, and virtualization strategy. You can choose a predefined profile suitable for the application that is using the storage, or you can create a custom profile.
storage resource card (SRC)	A card that provides the storage processors for the Data Services Platform (DSP). An SRC is always paired with a storage I/O (SIO) card. See also storage I/O (SIO) card.
storage resource card (SRC) set	Two cards in the Data Services Platform (DSP) that together provide the Fibre Channel (FC) and Gigabit Ethernet interfaces: the storage resource card (SRC) and the storage I/O (SIO) card. Two to four SRC sets can be installed in the DSP to provide the ports for connecting data hosts to the system. See also storage I/O (SIO) card.
Storage Service Processor (SSP)	A management device integrated into the system that provides unified management access to system components, as well as local and remote management and maintenance. The Storage Service Processor also supports automatic upgrades of patches, firmware, and software.

Storage Service Processor (SSP)

accessory tray The portion of the Storage Service Processor that contains a modem, router with a firewall, Ethernet hub, and network terminal concentrator (NTC).

storage tray An enclosure containing disks. A storage tray with a RAID controller is called a controller tray; a storage tray without a controller is called an expansion tray. See also controller tray and expansion tray.

store-and-forward See asynchronous replication

stripe size The number of blocks in a stripe. A striped array's stripe size is the stripe depth multiplied by the number of member extents. A parity RAID array's stripe size is the stripe depth multiplied by one less than the number of member extents. (SNIA) See also striping.

striping Short for data striping; also known as RAID Level 0 or RAID 0. A mapping technique in which fixed-size consecutive ranges of virtual disk data addresses are mapped to successive array members in a cyclic pattern. (SNIA) See also concatenation.

SunMC See Sun Management Center (SunMC).

Sun Management Center (SunMC)

An element management system for monitoring and managing the Sun environment. Sun Management Center also integrates with the leading enterprise management software, including the Storage Automated Diagnostic Environment, to provide customers with a unified management infrastructure. The base package of Sun Management Center is free and provides hardware monitoring. Advanced applications (add-ons) extend the monitoring capability of the base package.

Sun StorEdge Remote Response service

A remote support solution for Sun StorEdge series storage systems that proactively identifies operational anomalies to help prevent them from becoming business problems. Through around-the-clock monitoring, connectivity with Sun, and remote support, the Sun StorEdge Remote Response service helps maximize a storage system's availability.

suspend operation In the context of data replication, an operation in which replication set or consistency group activity is temporarily stopped, and an internal bitmap tracks write operations to the volume rather than sending the write operations over the physical link to the secondary volume. This method tracks write operations that have not been remotely copied while access to the secondary peer is interrupted or impaired. The software uses this replication bitmap to reestablish data replication through an optimized update synchronization rather than through a complete volume-to-volume copy. See also Fast Start operation and resume operation.

switch fabric card (SFC)	A board that provides the central switching function for the Data Services Platform (DSP).
synchronization	The act of aligning or making entries be equivalent at a specified point in time. (SNIA) In the context of mirroring, see resilvering.
synchronous replication	A replication technique in which data must be committed to storage at both the primary site and the secondary site before a write to the primary volume is acknowledged. See also asynchronous replication.
target	The system component that receives a SCSI I/O command. (SNIA) A target in the Sun StorEdge 6920 system can be an initiator or logical unit number (LUN).
thin-scripting client	See remote scripting CLI client.
topology	A graphical depiction of a storage network or storage system.
tray	See storage tray.
USB flash disk	A disk connected to the Storage Service Processor that stores system characteristics for the Storage Service Processor and Data Services Platform (DSP). The disk connects to the USB port on the Storage Service Processor.
virtual disk	A set of disk blocks presented to an operating environment as a range of consecutively numbered logical blocks with disk-like storage and I/O semantics. The virtual disk is the disk array object that most closely resembles a physical disk from the operating environment's viewpoint. (SNIA) In the Sun StorEdge 6920 system, the system itself is the operating environment.
virtualization	The act of integrating one or more (back end) services or functions with additional (front end) functionality for the purpose of providing useful abstractions. Typically virtualization hides some of the back end complexity, or adds or integrates new functionality with existing back end services. Examples of virtualization are the aggregation of multiple instances of a service into one virtualized service, or to add security to an otherwise insecure service. Virtualization can be nested or applied to multiple layers of a system. (SNIA) The Sun StorEdge 6920 system uses virtualization attributes to create and manage storage pools. See also concatenation and striping.
virtualization strategy	Selection of the technique used when virtualizing data across multiple virtual disks. See also concatenation and striping.

volume A logically contiguous range of storage blocks allocated from a single pool and presented by a disk array as a logical unit number (LUN). A volume can span the physical devices that constitute the array, or it can be wholly contained within a single physical disk, depending on its virtualization strategy, size, and the internal array configuration. The array controller makes these details transparent to applications running on the attached server system.

World Wide Name

(WWN) A unique identifier for a port, initiator, virtual disk, or volume, assigned by the system. The WWN of an object does not change throughout its lifetime and is never reused to name another object.

write order

consistency Preservation of write ordering across all volumes in a consistency group or in replication sets.

write ordering The process by which write operations that are directed to the secondary volume occur in the same order as write operations to the primary volume.

WWN See World Wide Name (WWN).

zone A collection of Fibre Channel N_Ports and/or NL_Ports (that is, device ports) that are permitted to communicate with each other via the fabric. (SNIA)

zoning A method of subdividing a storage area network (SAN) into disjoint zones, or subsets of nodes on the network. SAN nodes outside a zone are invisible to nodes within the zone. With switched SANs, traffic within each zone can be physically isolated from traffic outside the zone. (SNIA) See also zone.

Index

A

account passwords, changing 65

acknowledged alarms
deleting 211
reopening 210

actionable events 31

admin user name 18

Administration tab

for configuration administration 37
for monitoring administration 37

administrative functions

configuration 37
for monitoring administration 37

agent

activating and deactivating 78
overview of 28
running manually 79

agent statistics

clearing 79
displaying 79

aggregated events 31

displaying information about 208

alarm notification, types of 71

alarms

acknowledging 210
clearing 29
current 36
deleting 211
described 29
displaying detailed information about 208
displaying for a specific device 209
displaying statistics reports for 209

displaying summary information about 208
fault signature analysis of 69
generating statistics reports for 211
overview of 207
reopening 210
sorting 40
types of 36

Alarms tab 37

alternate master controller trays, location of 186

archived monitoring jobs

deleting 235
displaying a summary of 233
displaying details for 233

archived topologies

comparing snapshots of 227
deleting snapshots of 227
displaying 226
displaying snapshots of 226

array capacity, addition and removal of 240

array configurations 12

array controllers, addition and removal of 240

array hot-spares 4, 159

using disks as 151, 185

array type parameter, for storage profiles 160

assets reports, display settings for 206

asterisk, meaning of

in forms 41
in searches 42

asynchronous data replication

overview of 117
planning 130

asynchronous queue 129

- asynchronous replication
 - license for 83
- asynchronous replication mode 129
 - versus synchronous mode 131
- autosynchronization option 126
- auxiliary port 6
- B**
- backout reports, generating 223
- backup tape
 - data synchronization and 148
 - Fast Start suspend replication option and 133
- banner, described 34
- base cabinet
 - FRUs in 240
 - hardware components of 5
- base cabinets
 - numbering scheme for 186
- blocking mode, asynchronous replication and 129
- break operation
 - described 94
 - performing 101
- browser interface 14
 - navigating 34
 - page layout in 34

- C**
- cabinets, addition of 240
- CD-ROM, ejecting 80
- chassis LEDs 243
- column sort order, changing 40
- command-line interface, overview of 58
- concatenation
 - described 161
 - storage profile setting for 161
- configuration
 - default 20
 - exporting 70
- configuration file, for device inventory 28
- configuration jobs
 - accessing 37
 - canceling and removing 231
 - displaying information about 230
 - overview of 228
 - viewing 35
- configuration software 14

- configuration tasks, overview of 26
- consistency groups
 - adding replication sets to 142
 - asynchronous replication and 130
 - changing the properties of 145
 - changing the replication mode of 146
 - combining replication sets in 143
 - creating 141
 - deleting 145
 - displaying information about 140
 - overview of 121
 - planning 122, 137
 - removing replication sets from 144
 - replication sets and 117
 - resume replication operations and 133
- console, accessing 34
- content pane, elements in 38
- controller trays 4
 - configuration of, in arrays 181
 - described 184
 - location of 185
 - master 186
 - maximum number of 187
- cooling fans, for power supplies 11
- core files
 - downloading and deleting 258
 - overview of 258
- cron daemon
 - agent and 29
- current alarms 36
- current jobs, canceling 231
- current user logins 36

- D**
- data hosts
 - expanding the number of 56
 - FC ports and 175
 - multipathing and 57
 - overview of 51
 - ports for 54
 - service panel ports for 53
 - storage domains and 167
- data loss, replication mode and 132
- data migration, from external storage 201
- data migration, mirroring and 93
- data partitions, in virtual disks 151
- data paths, external storage and 199

- data replication
 - autosynchronization option in 126
 - changing modes in 146
 - consistency groups in 121
 - disabling ports for 178
 - enabling ports for 177
 - operation options in 132
 - overview of 116
 - peer role reversal in 145
 - performing 147
 - planning 135
 - port requirements for 175
 - replication links in 123
 - replication modes in 129
 - replication set properties in 126
 - replication sets in 118
 - reversing roles in 128
 - suspended mode and 134
 - task overview for 136
 - troubleshooting 148
 - write order consistency in 121
- data replication operations, planning 132
- data services
 - about licenses for 82
 - adding license keys for 84
 - deleting license keys for 85
 - displaying license key information for 85
 - overview of 17
- Data Services Platform (DSP)
 - displaying performance data for 217
 - displaying utilization and replication queue information for 217
 - firmware upgrades to 238
 - FRUs in 240
 - LEDs on 242
 - overview of 9
 - reserved IP address for 47
 - reserved LAN name for 47
 - troubleshooting 242
- data set
 - snapshot policy and 111
- data sets
 - independent versus dependent copies of 99
 - mirroring and 93
- data striping
 - described 161
 - and RAID levels 150
 - storage profile setting for 161
- data synchronization
 - auto versus manual 126
 - displaying jobs for 133
 - overview of, in mirroring 93
 - using a backup tape for 148
- dedicated hot-spares 4, 159
 - setting for, in storage profiles 159
- default configuration 20
 - changing 26
- default storage configuration 21
- default storage domain 167
 - initiators and 171
 - restrictions on deleting 170
- Default storage pool
 - DEFAULT storage domain and 20
- Default storage profile 162
 - characteristics of 19
 - restrictions on modifying or deleting 20
- dependent data set copies 99
- device inventory, configuration file for 28
- devices
 - adding notes to displayed information for 218
 - changing the monitoring status of 219
 - changing the name of 219
 - discovering 37
 - displaying alarms for 209
 - displaying detailed information about 214
 - displaying FRU or component inventory listings for 214
 - displaying system parameters for 216
 - enabling and disabling monitoring of 219
 - installing patches for 222
 - running diagnostic tests on 259
 - updating monitoring and setup information for 219
 - updating the inventory list for 219
 - viewing FRU or component details about 215
- DHCP port 52
- diagnostic software 14
 - external storage and 199
- diagnostic software, described 28
- diagnostics
 - running 259
- disaster recovery strategies, replication sets and 25
- discovery events 30
- disk drive setting, for storage profiles 160

- disk drives 25
 - addition and removal of 240
 - associating with storage profiles 160
 - characteristics of 189
 - displaying 190
 - maximum and minimum number of 189
 - maximum number of, per array 180
 - online removal of 189
 - and trays 12
- disk scrubbing, Volume Verify and 220
- display filters, applying 39
- DNS port 52
- drive capacity 189
- drive sizes 189

E

- ejecting a CD-ROM 80
- email filters
 - adding 75
 - deleting 75
 - editing 75
 - overview of 74
- email notification
 - adding a recipient for 73
 - configuring 72
 - deleting a recipient for 74
 - described 71
 - editing a recipient for 73
 - testing addresses for 72, 256
 - using filters for 74
- error conditions, indications of 207
- error messages, listing of 260
- Ethernet connections
 - for expansion cabinets 6
- Ethernet hub 9
- Ethernet ports, for MICs 11
- ethers
 - entering and editing 70
- event generation, overview of 30
- event life-cycle, described 30
- event log
 - displaying 212
- event maximums database 30
- events
 - actionable 31
 - aggregated 31

- aggregated, displaying information about 208
 - causes of 29
 - displaying detailed information about 212
 - displaying the log for 212
 - email notification for 74
 - fault signature analysis of 69
 - filtering display of 39
 - overview of 28, 207
- events logs 30
- expansion cabinets
 - cabling 7
 - cabling for 7
 - FRUs in 240
 - in numbering scheme for 186
 - number of arrays in 180
- expansion trays
 - described 184
- external modems
 - addition and removal of 240
- external storage 25
 - connecting 192
 - displaying information about 199
 - legacy volumes and 194
 - migrating data from 201
 - planning 194
 - troubleshooting 199
 - using as raw storage 202
- external storage devices
 - configuration of 193
 - depiction of 191
 - ports and 175
 - storage domains and 24
 - supported 193
 - viewing 37
- External Storage tab 37
- external switches 51
 - connecting 56

F

- failover, and controller trays 186
- fan assembly 11
- fan LEDs 244
- Fast Start suspend operation, in data replication 133
- fault signature analysis
 - enabling and disabling 69
- fault tolerance, and RAID levels 150
- feature license keys 82

- adding 84
- deleting 85
- displaying 85
- feature licenses
 - displaying compliance information about 84
 - license keys and 82
 - overview of 82
 - transferring 85
- Fibre Channel (FC) connections
 - for expansion cabinets 6
- Fibre Channel (FC) ports 175
 - arrays and 180
 - data host connections and 51
 - disabling for data replication 178
 - displaying information about 176
 - enabling for data replication 177
 - guidelines for 55
 - initiators and 171
 - replication links and 123
- field-replaceable units (FRUs)
 - displaying details about 215
 - displaying inventory information about 214
 - replacement of 240
 - servicing 239
- filters
 - using for email notification 74
- filters, display, applying 39
- firewall, limitations on 47
- firewall, system 9
- firmware
 - upgrading 15
- firmware upgrades 238
 - backing out of 238
- flash disk, resident 9
- force break operation
 - described 94
 - performing 102
- Full resume synchronization option 132
- full shutdown 59
 - performing 60
 - restoring the system after 61

G

- general setup parameters
 - displaying 68
 - editing 68
- Gigabit Ethernet connections 6

- Gigabit Ethernet ports 123, 175
 - disabling for data replication 178
 - displaying information about 176
 - displaying replication link information about 139
 - enabling for data replication 177
 - replication links and 123
- guest user name 18

H

- hardware components
 - displaying details about 215
 - displaying inventory information about 214
 - servicing 239
- Help button 35
- help function, using 42
- high availability
 - tray configuration for 180
- high availability, effect of replication mode on 132
- High Performance Computing storage profile 162
- high-availability (HA) features 4
- historical jobs, removing 231
- host bus adapters (HBAs) 24
 - data host connections and 51
 - initiators and 171
 - overview of 55
- host names 48
 - creating 49
 - for arrays 180
 - for internal routers 48
 - for network terminal concentrator (NTC) 48
 - for storage arrays 47
 - for Storage Service Processor 48
- host software
 - downloading 16
 - for Solaris OS 15
 - multipathing 16
- hot-spare options, for storage profiles 159
- hot-spares
 - setting for, in storage profiles 159
 - size of 189
 - types of 4, 189
- HTTP service, monitoring software and 28

I

- I/O panel 53
- I/O performance, versus synchronization rate 138

- independent data set copies 99
- Initiator Delete job 229
- Initiator Mapping job 229
- initiators 24
 - creating 172
 - data host connections and 51
 - deleting 173
 - deleting volumes associated with 90
 - displaying information about 172
 - host bus adapters (HBAs) and 55
 - jobs for operations on 229
 - mapping a volume to 90
 - mapping legacy volumes to 194
 - mapping to volumes 173
 - number of 54
 - overview of 171
 - preconfiguring 172
 - restrictions on deleting 173
 - in a SCSI environment 25
 - unmapping a volume from 91
 - unmapping from a volume 174
 - volumes and 87
 - World Wide Name (WWN) for 172
- interface, navigating 34
- Inventory tab 37
- inventory, generating 219
- IP addresses
 - creating 49
 - for Data Services Platform (DSP) 47
 - MICs and 4
 - setting 63
 - for arrays 180
 - for storage arrays 47
- isolated system 57
- isolation policy, for mirror components 95

J

- Java Web Console, accessing 34
- Job Cancel job 229
- Job Remove job 229
- jobs
 - configuration
 - canceling and removing 231
 - displaying information about 230
 - overview of 228
 - monitoring
 - about 232

- archived, deleting 235
 - archiving 234
 - deleting 235
 - displaying a summary of 232
 - displaying an archive summary of 233
 - displaying archive details for 233
 - displaying details about 233
 - stopping 234
 - viewing 35
- Jobs tab
 - for configuration jobs 37
 - for monitoring jobs 37

L

- LAN connections 6
- LAN names, predefined 47
- LEDs
 - for cards 245
 - chassis 243
 - on Data Services Platform (DSP) 242
 - fan 244
 - overview of 6, 241
 - power supply 244
- legacy volume
 - limitations on replicating 92
- legacy volumes
 - adding to storage pools 197
 - creating 200
 - described 194
 - mirroring and 93
 - replicating data to and from 120
 - restrictions on extended 120
 - troubleshooting 199
 - versus raw storage 197
- licensing requirements, complying with 82
- lights-out management
 - addition, removal, and disabling of 240
 - overview of 58
- log events, generation of 30
- log files
 - events and 29
 - viewing 271
- Log Out button 35
- logging out 43
- logical storage elements 25
 - configuring 37
 - physical storage elements and 21

- searching for 41
- Logical Storage tab 37
- logins
 - current user 36
- logins, user 18
- loop cards (LPCs) 4
- LUN masking 171
- LUN requirements, for external storage devices 194
- luxadm(1M) command 55

M

- Mail Spool storage profile 162
- management host
 - and isolated system 57
 - thin-scripting client and 58
- management interface cards (MICs) 4
 - DSP 11
 - IP addresses and 47
 - LEDs for 245
- management software 14
- manual synchronization 126
- Mapping Delete job 229
- master controller trays, location of 186
- master/alternate master modes 4
- media access control (MAC) addresses
 - entering and editing 70
- mirror components
 - adding 99
 - breaking 101
 - described 25, 93
 - displaying information about 97
 - extension of volume capacity and 91
 - forcing the break of 102
 - planning 95
 - rejoining 100
 - reverse rejoining 101
 - splitting 99
 - versus creating snapshots 96
- mirrored volumes
 - capacity of 95
 - creating from an existing volume 98
 - creating new 97
 - described 93
 - expanding the snapshot reserve space for 109
 - planning 95
 - snapshot rollbacks and 105

- mirroring
 - external data migration and 201
 - legacy volumes and 194
 - overview of 93
 - predefined storage profile for 163 and RAID levels 150
- mirrors
 - adding components to 99
 - described 25, 93
 - methods for creating 95
- modem 9
- modem connection 6
- monitoring cycle
 - described 29
- monitoring jobs
 - about 232
 - accessing 37
 - archived, deleting 235
 - archiving 234
 - deleting 235
 - displaying a summary of 232
 - displaying an archive summary 233
 - displaying archive details for 233
 - displaying details about 233
 - stopping 234
- monitoring software 14
- monitoring software, described 28
- monitoring status
 - changing for a device 219
- monitoring strategy, overview of 29
- monitoring tasks, overview of 32
- multipathing software
 - for Solaris OS 15
 - for third-party platforms 16
- multipathing, overview of 57

N

- navigational tabs 36
- network configuration
 - options for 50
 - site LAN and 51
- network domains, types of 46
- network response time, effect of replication mode on 131
- network terminal concentrators (NTCs) 9
- NFS Mirroring storage profile 163

- NFS Striping storage profile 163
- Normal resume synchronization option 132
- Normal suspend operation, in data replication 133
- notification setup parameters, editing 78
- notification, types of 71
- NSCC notification, described 71
- NTP server, enabling 64
- number of disks, setting in storage profiles 160

O

- operating system upgrades, backing out of 238
- Oracle DSS storage profile 163
- Oracle OLTP HA storage profile 163
- Oracle OLTP storage profile 163

P

- packet processing 11
- page banner, described 34
- page display controls 39
- parent volumes
 - described 103
 - rolling back a snapshot for 114
 - setting the snapshot policy for 111
 - snapshot reserve space and 104
 - snapshot rollbacks and 105
 - write operations and 104
- parity checking, and RAID levels 150
- partial shutdown
 - described 59
 - performing 59
 - restoring the system from 60
- partitions, in virtual disks 151
- passwords
 - changing 65
 - default 18
- patch upgrades 238
 - backing out of 238
- patches
 - about 221
 - backing out 224
 - installing 222
- PCMCIA device 11
- peer roles
 - consistency group requirements for 122
 - described 118

- reversal of 118, 128
- performance
 - and RAID levels 150
 - system, storage allocation and 23
- performance data
 - displaying for DSP or array 217
 - displaying for the rack 216
- physical storage elements 24
 - configuring 37
 - logical storage elements and 21
 - searching for 41
- Physical Storage tab 37
- Pool Delete job 229
- ports
 - allocating 53
 - disabling for data replication 178
 - displaying information about 176
 - displaying replication link information for 177
 - enabling for data replication 177
 - legacy volumes and 198
 - order of, for storage connections 54
 - overview of 25, 175
 - replication sets and 120
 - SRC set 54
- power connections
 - for expansion cabinets 6
- power cords 11
- power relay
 - partial shutdown and 59
- power supply
 - multiple 11
 - overview of 9
- power supply LEDs 244
- powering off 59
 - full shutdown 60
 - partial shutdown 59
- powering on 61
- preference settings, for diagnostic topologies and asset reports 206
- Preferences button 35
- pre-installed software 14
- primary peers
 - consistency group requirements for 122
 - effect of role reversal on 128
 - replication links and 123
 - reversing the role of 145

- primary volumes
 - autosynchronization and 126
 - consistency groups and 122
 - replicating data from 147
 - replication modes and 129
 - replication sets and 117
 - resume replication options and 132
- Profile Delete job 229
- provisioning storage 21

R

- racks
 - adding notes to displayed information for 218
 - displaying components of 214
 - displaying inventory of 213
 - displaying performance data for 216
- RAID controllers 184
- RAID levels, supported 150
- RAID management hardware 13
- RAID sets 150
 - described 25
- RAID systems, overview of 150
- Random 1 storage profile 162
- Random 5 storage profile 162
- raw storage
 - adding to storage pools 202
 - pool requirements for 197
 - versus legacy volumes 197
- read requests, snapshots and 104
- read-ahead mode, for storage profiles 160
- Refresh button 34
- rejoin operation
 - described 94
 - performing 100
- remote notification
 - configuring 76
 - types of 71
- remote power management 58
- remote scripting CLI client 14
- remote scripting client, installing 58
- replicated volumes
 - extending 92
 - managing 120
 - overview of 117
 - snapshot rollbacks and 105
- replication bitmaps

- described 119
- Normal replication suspend operation and 133
- replication sets and 117

- replication links
 - adding 179
 - deleting 179
 - displaying information about 139, 177
 - overview of 123
 - resume replication operations and 133
- replication modes
 - changing 146
 - described 129
 - planning 131
 - restrictions on changing 131
- replication queue depth, displaying 217
- Replication Set Delete job 229
- replication sets
 - adding to a consistency group 142
 - asynchronous replication and 130
 - changing the properties of 145
 - changing the replication mode of 146
 - combining in a consistency group 143
 - consistency groups and 121
 - creating 141
 - deleting 142
 - deleting from consistency groups 144
 - described 25
 - displaying information about 140
 - effect of consistency group operations on 122
 - effect of deleting 120
 - overview of 118
 - planning 120, 137
 - replication modes for 129
 - requirements for 121
 - resume replication operations and 133
 - setting the properties of 126
 - suspended mode and 134
- Replication Synchronize job 230
- Reset button 41
- resident software 14
- resilver priority, deciding 96
- Resilvering Mirror Components job 230
- resilvering process
 - overview of 93
 - rejoin operation and 100
- resume replication options, described 132
- reverse rejoin operation

- described 94
- performing 101
- revision analysis
 - generating backout reports for 223
 - generating upgrade reports for 222
- revision maintenance
 - generating backout reports for 223
 - generating upgrade reports for 222
- revision maintenance software 15
- revision patches
 - about 221
 - backing out 224
 - installing 222
- role reversal
 - overview of 128
 - performing 145
- roles, user 18
- Rollback Volume job 230
- router, system 9

S

- SAN Foundation software
 - for Solaris OS 15
- Save button 41
- SCSI adapters 25
- SCSI devices, initiators and 171
- search function 35
 - in help 43
 - using 41
- secondary peers
 - consistency group requirements for 122
 - effect of role reversal on 128
 - replication links and 123
 - reversing the role of 145
- secondary volumes
 - autosynchronization and 126
 - consistency groups and 122
 - methods for updating 119
 - replicating data to 147
 - replication modes and 129
 - replication sets and 117
 - resume replication options and 132
 - snapshot rollbacks and 128
 - snapshots of 128
- Secure HTTP service, monitoring software and 28
- security considerations 50

- in storage allocation 22
- segment size parameter, for storage profiles 159
- Sequential storage profile 162
- serial ports 6
 - for MICs 11
- Service Advisor 35
 - accessing 239
 - button for 35
 - FRU replacement and 240
- service panel
 - described 53
 - expansion cabinet cabling and 7
 - overview of 6
- Service Processor panel 53
- services, restricting 52
- session time, jobs and 228
- shutdown, types of 59
- SIO Combo cards, LEDs for 254
- site information
 - displaying 67
 - editing 67
- site LAN 46
 - connecting to the system 51
 - isolating the system from 57
 - managing the system from 51
 - restricting access from 52
 - system security and 50
- SMTP port 52
- snapshot count, specifying 105
- Snapshot Delete job 229
- snapshot policy
 - rollbacks and 106
 - selecting 111
- snapshot reserve space
 - allocating 108
 - changing the allocation of 107
 - deleting 110
 - expanding 109
 - freeing 113
 - overview of 104
 - planning 106
 - snapshot rollbacks and 105
- Snapshot Resnap job 229
- snapshot rollbacks
 - canceling 114
 - overview of 105

- performing 114
- secondary volumes and 128
- snapshot policy and 112
- snapshots
 - automatic creation option for 127
 - compared with split mirror components 96
 - creating 112
 - deleting 113
 - deleting volumes associated with 90
 - displaying information about 111
 - extension of volume capacity and 91
 - of legacy volumes 194
 - license for 83
 - managing 111
 - maximum number of 105
 - mirror components and 99
 - overview of 103
 - planning 106
 - replacing 113
 - in a replication set 120
 - of replication sets 122
 - retaking 113
 - rolling back 114
 - selecting the snapshot policy for 111
 - specifying write activity for 105
 - synchronization and 127
- SNMP notification
 - adding a recipient for 76
 - deleting a recipient for 77
 - editing a recipient for 77
- SNMP traps, described 71
- software upgrades 238
- software version, displaying 34
- solution extract
 - displaying 257
 - generating 257
- split mirror components
 - compared with snapshots 96
 - creating 99
 - rejoining 100
 - reverse rejoining 101
- split operation
 - described 94
 - performing 99
- sscs command 58
- statistics reports
 - displaying 209
 - generating 211
- storage area networks (SANs)
 - initiators in 171
- storage arrays 25
 - configuration of 180
 - displaying information about 182
 - displaying performance data for 217
 - host names and IP addresses for 180
 - hot-spares in 189
 - modifying 183
 - number of 54, 180
 - overview of 12
 - ports for 53
 - read-ahead mode for 160
 - replacement of 240
 - reserved names and IP addresses for 47
 - segment size in 159
 - supported configurations for 12
 - tray configuration in 184
 - typical configuration for 181
- storage configuration, changing 26
- storage customers, authentication of 171
- storage domains
 - configuration options for 20
 - creating 169
 - deleting 170
 - displaying information about 169
 - external storage devices and 24
 - initiators in 171
 - maximum number of 168
 - mirrors and 95
 - overview of 167
 - replication sets and 121
 - storage pools and 153
 - storage profiles and 158
- storage elements
 - logical 25
 - physical 24
- storage I/O (SIO) cards 11
 - LEDs for 252
- storage I/O (SIO) Combo cards 54
- storage pools 25
 - adding external storage to 197
 - adding storage to 155
 - configuration options for 20
 - creating 155
 - default 21
 - deleting 156

- displaying information for 154
 - jobs for operations on 229
 - legacy volumes and 194
 - modifying 156
 - multipathing and 57
 - overview of 153
 - planning 153
 - read-ahead mode requirements for 160
 - reconfiguring 154
 - restrictions on deleting 154
 - restrictions on deleting profiles associated with 20
 - storage domains and 168
 - virtual disks and 151
 - volumes and 87
- storage profile parameters, described 158
- storage profiles
- choosing a virtualization strategy for 161
 - configuration options for 20
 - creating 165
 - Default 19
 - deleting 166
 - displaying information about 164
 - jobs for operations on 229
 - modifying 165
 - optimizing attributes with 23
 - overview of 158
 - predefined 162
 - restrictions on modifying or deleting 20
 - storage pools and 153
- storage resource card (SRC) sets 11
- overview of 54
 - ports and 175
 - replication sets and 121
- storage resource cards (SRCs) 11
- LEDs for 251
 - removal and addition of 240
- Storage Service Processor
- internal LAN name and internal IP address for 48
 - overview of 8
 - reserved IP address for 47
 - reserved LAN name for 47
 - upgrades to 238
- Storage Service Processor (SSP) LAN 46
- Storage Service Processor accessory tray
- overview of 9
- storage trays 25
- configuration of 180, 184
 - displaying information about 188
 - maximum number of 187
 - number of 180
 - number of virtual disks in 151
 - numbering of 185
 - overview of 184
 - sizes of drives in 189
 - typical configuration for 181
- storage user name 18
- storage volumes
- associating an initiator with 172
 - changing the name or description of 90
 - creating 89
 - deleting 90
 - described 25, 87
 - displaying information for 89
 - effect of consistency group operations on 122
 - extended the capacity of 91
 - extending 120
 - external 194
 - jobs for operations on 229
 - legacy 194
 - mapping an initiator to 173
 - mapping to initiators 90
 - maximum number of 198
 - mirroring 98
 - planning 88
 - replicated 117
 - replicated, extending 92
 - replication sets and 121
 - requirements for 88
 - restrictions on deleting 90
 - restrictions on mirroring 120
 - standalone, from mirrors 101
 - storage pools and 154
 - unmapping an initiator from 174
 - unmapping from initiators 91
- storage, adding to a pool 155
- stripe size, setting for storage profiles 161
- Sun StorEdge Data Mirror software 17
- license for 82
- Sun StorEdge Data Replicator software 17
- license for 83
- Sun StorEdge Data Snapshot software 17
- license for 83
- Sun StorEdge Remote Response (SSRR) service
- component updates and 238

- modem for 9
- system IP addresses and 49
- Sun StorEdge Remote Response service
 - described 71
- Sun StorEdge Remote Response service software
 - overview of 15
- Sun StorEdge SAN Foundation software
 - external switches and 56
 - host bus adapters and 55
- Sun StorEdge Storage Pool Manager software 17
 - license for 83
- Sun StorEdge Traffic Manager software 57
- supported RAID levels 150
- suspend replication operation
 - options for 133
 - planning 134
- suspended mode
 - asynchronous replication and 129
 - autosynchronization and 126
- switch fabric cards (SFCs) 12
 - LEDs for 249
- switches
 - external 51
 - external, connecting 56
- Sybase DSS storage profile 164
- Sybase OLTP HA storage profile 164
- Sybase OLTP storage profile 163
- synchronization operations
 - options for 132
 - restarting after suspension 133
- synchronization rate, versus I/O performance 138
- synchronous data replication, overview of 117
- synchronous replication
 - license for 83
- synchronous replication mode 129
 - versus asynchronous mode 131
- syslog file 29
- system
 - restoring from a full shutdown 61
 - restoring from a partial shutdown 60
- system configuration
 - exporting 70
 - generating a snapshot of 257
- system defaults 19
- system description, adding 63
- system error messages, listing of 260

- system firewall 9
- system IP address, setting 63
- system messages
 - error, listing of 260
- system monitoring software, described 28
- system name, displayed in banner 35
- system overview 2
- system performance, storage allocation and 23
- system racks
 - displaying components of 214
 - displaying inventory of 213
- system router 9
- system security, options for 50
- system shutdown, types of 59
- system time
 - setting manually 65
 - setting with the NTP server 64
- system timeout settings
 - displaying 68
 - editing 68
- system topology, displaying 225

T

- table columns, changing the sort order of 40
- tables, filtering information in 39
- tabs, navigational 36
- threshold notifications, for snapshot reserve space 105
- topologies
 - archived, displaying 226
 - archiving 226
 - changing the display of 225
 - comparing archived snapshots of 227
 - deleting archived snapshots of 227
 - displaying archived snapshots of 226
 - displaying system inventory in 225
- topology
 - display settings for 206
- topology snapshots
 - comparing 227
 - deleting 227
 - displaying 226
- Topology tab 37
- topology views, changing 225
- trays. *See* storage trays

- U**
 - upgrade process, troubleshooting 238
 - upgrade reports, generating 222
 - USB flash disk
 - license keys and 84
 - USB flash disk connection 9
 - user name, displayed in banner 35
 - user roles 18
- V**
 - VERITAS Volume Manager with DMP 57
 - Version button 34
 - Virtual Disk Create job 230
 - Virtual Disk Delete job 230
 - virtual disks
 - configuration options for 20
 - deleting 151
 - described 25
 - display of, for external storage 197
 - displaying information about 151
 - external storage and 193
 - hot-spares and 159
 - hot-spares for 189
 - initialization of, for legacy volumes 196
 - method of creating 154
 - number of disk drives in 160
 - number of in tray 185
 - overview of 150
 - read-ahead mode requirements for 160
 - segment size requirements for 159
 - snapshot reserve space and 104
 - in storage pool planning 153
 - volumes and 87
 - virtualization
 - of storage assets 17
 - virtualization attributes, selecting 161
 - virtualization services 9
 - virtualization strategy
 - choosing 161
 - setting for, in storage profiles 161
 - volume capacity
 - legacy volumes and 194
 - volume capacity, extending 91
 - Volume Delete job 229
 - Volume Map job 229
 - Volume Snapshot job 229
 - volume snapshots
 - described 25
 - jobs for operations on 229
 - Volume Verify function
 - setting up 220
- W**
 - web browser
 - content pane elements in 38
 - navigating 34
 - system IP address and 64
 - wildcard character, in searches 42
 - write activity
 - snapshot reserve space and 105
 - write operations
 - parent volume and 104
 - write operations, effect of replication mode on 131
 - write order consistency
 - consistency groups and 121
 - described 117
 - effect of replication mode on 132
 - write response time, effect of replication mode on 132
 - WWN (World Wide Name)
 - for initiators 171
- Z**
 - zoning, external storage and 195