



Sun StorageTek™ Array Administration Guide

For the Browser Interface Management Software

Sun Microsystems, Inc.
www.sun.com

Part No. 819-5050-10
May, 2006

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

Copyright 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

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

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

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

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

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

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

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

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

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

Copyright 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

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

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

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

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

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

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

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

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



Please
Recycle



Adobe PostScript

Contents

Preface xv

1. Overview 1

Array and Software Overview 2

 About the Sun StorageTek 6140, 6130, and 6540 Arrays 2

 About the Array Software 3

Hardware Components Overview 7

 About High-Availability Features 7

 About the Sun StorageTek 6140 Storage Array Hardware 8

 Hardware Overview 8

 About the Sun StorageTek 6140 Controller Tray 9

 About the Sun StorageTek 6140 Expansion Tray 12

 About the Sun StorageTek 6130 Storage Array Hardware 13

 Hardware Overview 13

 About the Sun StorageTek 6130 Controller Tray 14

 About the Sun StorageTek 6130 Expansion Tray 15

2. Using the Browser Interface 17

 About the Browser Interface 18

 About the Page Banner 18

About the Navigation Tree	19
About the Page Content Area	20
Controlling the Display of Table Information	21
About the Status Icons	22
Using Forms	23
Searching for System Elements	24
Using Help	25
Logging Out of the Management Software	26

3. Performing Connection and Administration Tasks 27

Connection Tasks	28
Managing From the Site LAN	28
Managing an Array Isolated From the LAN	29
Installing the Remote CLI Client	29
About Host Bus Adapters	30
About Multipathing	30
Administration Tasks	31
Accessing, Adding, and Deleting Storage Arrays	31
Displaying Array Information	31
Planning to Register an Array	32
Registering an Array	32
Unregistering an Array	33
Upgrading Array Firmware	33
Monitoring Array Health	33
Specifying General Settings	34
About the Array Password	34
Changing the Array Password	35
Setting Array Details	36
Enabling Disk Scrubbing	36

Setting the Time	38
Setting the Array IP Address	38
Managing User Accounts	39
About User Accounts	40
Displaying User Information	40
Adding a New User	41
Removing a User	41
Managing Licenses	42
About Licensed Features	42
Displaying License Information	42
Adding a License	43
Disabling a License	43
Re-enabling a License	44
Displaying Event Information	44
Monitoring Performance	45
Administering Jobs	45
About Jobs	45
Displaying Job Information	46
Canceling Jobs	46
Viewing the Activity Log	46
4. Performing Configuration Tasks	49
Initial State of the Array	50
About the Initial State of the Array	50
About Provisioning Storage	50
About Physical Storage Elements	51
About Logical Storage Elements	52
Configuring Storage Volumes	53
About Volumes	53

Planning Volumes	54
Managing Volumes	56
Displaying Volume Information	56
Displaying Volume Performance Statistics	57
Creating a Volume	57
Modifying a Volume	57
Copying Volume Information	58
Mapping a Volume to a Host or Host Group	58
Expanding Volume Capacity	59
Unmapping a Volume From a Host or Host Group	60
Changing Controller Ownership of a Volume	60
Deleting a Volume	61
Configuring Volume Copies	62
About Volume Copies	62
Planning Volume Copies	63
Managing Volume Copies	65
Enabling the Volume Copy Feature	65
Displaying Volume Copy Information	66
Creating a Volume Copy	66
Recopying a Volume Copy	67
Changing the Copy Priority	68
Removing a Copy Pair	69
Configuring Volume Snapshots	70
About Volume Snapshots	70
Planning Volume Snapshots	72
Calculating Reserve Volume Capacity	73
Managing Volume Snapshots	76
Enabling Volume Snapshots	76

Displaying Volume Snapshot Information	77
Creating a Volume Snapshot	77
Resnapping a Volume Snapshot	78
Copying a Volume Snapshot	78
Mapping a Volume Snapshot to a Host or Host Group	79
Expanding Snapshot Capacity	80
Unmapping a Volume Snapshot	80
Disabling a Volume Snapshot	81
Deleting a Volume Snapshot	82
Configuring Data Replication	83
About Data Replication	83
About the Replication Process	83
About Replication Sets	85
About the Consistency Group	87
About Replication Links	88
About Replication Set Properties	88
About Primary and Secondary Role Reversal	89
About Replication Modes	91
Reference: Synchronous Versus Asynchronous Replication Modes	92
About Data Replication Operations	93
Planning for Data Replication	93
General Planning Considerations	94
Planning to Create Replication Sets	95
Managing Data Replication	97
About Data Replication Status	97
Enabling Data Replication	98
Activating and Deactivating Data Replication	99
Disabling Data Replication	100

Displaying Replication Set Information	100
Creating Replication Sets	101
Deleting Replication Sets	102
Reversing Roles	102
Changing Replication Modes	103
Suspending and Resuming Data Replication	104
Testing Replication Set Links	105
Troubleshooting Data Replication	105
Configuring Virtual Disks	106
About Virtual Disks	106
Displaying Virtual Disk Information	107
Expanding a Virtual Disk	107
Defragmenting a Virtual Disk	108
Configuring Storage Pools	110
About Storage Pools	110
Displaying Pool Information	110
Creating a Storage Pool	111
Modifying a Storage Pool	112
Deleting a Storage Pool	112
Configuring Storage Profiles	113
About Storage Profiles	113
Reference: Standard Storage Profiles	114
Displaying Profile Information	115
Creating a Storage Profile	115
Modifying a Storage Profile	116
Copying a Storage Profile	116
Deleting a Storage Profile	117
Configuring Initiators	118

About Initiators	118
Displaying Initiator Information	118
Creating an Initiator	119
Deleting an Initiator	120
Configuring Host Groups and Hosts	121
About Host Groups	121
Managing Host Groups	122
Displaying Host Group Information	122
Creating a Host Group	122
Adding/Removing Host Group Members	123
Deleting a Host Group	124
Mapping a Host Group to a Volume	125
Unmapping a Host Group From a Volume	126
About Hosts	127
Managing Hosts	127
Displaying Host Information	127
Creating a Host	128
Deleting a Host	129
Mapping a Host to a Volume	129
Unmapping a Host From a Volume	130
Managing Trays and Disk Drives	132
About Trays and Disk Drives	132
Displaying General Controller Information	133
Displaying Controller Performance Statistics	133
Testing Controller Communication	134
Displaying Port Information	134
Displaying Tray Information	135
Displaying Disk Information	136

Managing Hot-Spare Drives	136
About Hot-Spare Drives	136
Planning Hot-Spares	137
Assigning and Unassigning Hot-Spares Manually	137
Assigning and Unassigning Hot-Spares Automatically	138
Configuring Storage Domains	139
About Storage Domains	139
Preparing to Create a Storage Domain	140
Determining How to Create a Storage Domain	140
Enabling the Storage Domain Feature	141
About Performing Mapping Functions	142
Displaying All Mappings on the Array	142
Creating a Mapping	143

5. Troubleshooting the System 145

Basic Troubleshooting	146
About Patches, Firmware, and Software	146
Backing Out Changes	147
About Field-Replaceable Units	147
About LED Status Indicators	147
Checking LEDs	156
Responding to an Array Boot-Up Failure	157
Viewing Log Files	157
Resetting a Storage Array Configuration	157
Redistributing Volumes	158
Resetting a Controller	159
Reviving Virtual Disks and Disk Drives	160

Glossary 161

Index 169

Preface

The *Sun StorageTek Array Administration Guide* describes how to configure the Sun StorageTek 6130, 6140 and 6540 arrays.

Before You Read This Book

Before you begin to configure the array, you must have already installed the array as described in these books:

- v *Sun StorageTek 6140 Array Getting Started Guide*
 - v *Sun StorageTek 6130 Array Getting Started Guide*
 - v *Sun StorageTek 6540 Array Release Notes*
-

How This Book Is Organized

Chapter 1 provides an overview of the array and management software.

Chapter 2 describes how to use the browser interface.

Chapter 3 describes connection and administration tasks, including general setup tasks.

Chapter 4 describes how to configure the array, including configuring volumes, volume copies, volume snapshots, storage pools and profiles, virtual disks, hosts and host groups, and storage domains.

Chapter 5 describes how to troubleshoot the array.

Using UNIX Commands

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

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

Shell Prompts

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

Typographic Conventions

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

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

Related Documentation

Application	Title	Part Number
Site planning information	<i>Sun StorageTek 6140 Array Site Preparation Guide</i>	819-5046- <i>nn</i>
Instructions for installing the array	<i>Sun StorageTek 6140 Array Getting Started Guide</i>	819-5045- <i>nn</i>
Late-breaking information not included in the information set	<i>Sun StorageTek 6140 Array Release Notes</i> <i>Sun Storage Automated Diagnostic Environment Enterprise Edition Release Notes</i>	819-5044- <i>nn</i> 819-0432- <i>nn</i>
Quick reference information for installing the array	<i>Sun StorageTek 6140 Poster</i>	819-5064- <i>nn</i>
Quick reference information for the CLI	<i>Sun StorageTek 6140 Array CLI Quick Reference Guide</i>	819-5051- <i>nn</i>
Regulatory and safety information	<i>Sun StorageTek 6140 Array Regulatory and Safety Compliance Manual</i>	819-5047- <i>nn</i>
Instructions for installing the Sun StorEdge Expansion cabinet	<i>Sun StorEdge Expansion Cabinet Installation and Service Manual</i>	805-3067- <i>nn</i>
Instructions for installing the Sun Rack 900/1000 cabinets	<i>Sun Rack Installation Guide</i>	816-6386- <i>nn</i>

Accessing Sun Documentation

You can obtain Sun network storage documentation at:

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

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

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

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 StorageTek Array Administration Guide, part number 819-5050-10

Overview

Click a link below for overview information about the Sun StorageTek arrays, including the 6140, 6130, and 6540:

- [“Array and Software Overview”](#) on page 2
- [“Hardware Components Overview”](#) on page 7

Array and Software Overview

This section describes the features of the array and user accounts:

- “About the Sun StorageTek 6140, 6130, and 6540 Arrays” on page 2
- “About the Array Software” on page 3

About the Sun StorageTek 6140, 6130, and 6540 Arrays

The Sun StorageTek arrays, including the 6130, 6140, and the 6540, are high-performance, enterprise-class, full 4 Gigabit per second (GB/s) for the 6140 and 6540 arrays, or 2 Gigabit per second (GB/s) for the 6130 array, Fibre Channel solutions that combine outstanding performance with the reliability, availability, flexibility, and manageability.

The 6130, 6140, and 6540 are modular, rackmountable and scalable arrays from a single controller tray (1x1) configuration to a maximum configuration of 1x8 with seven additional expansion trays behind one controller tray for the 6130 array, or ax7 with six additional expansion trays behind one controller tray for the 6140 array (FIGURE 1-1).

Refer to the Sun StorageTek 6540 Array Release Notes for information specific to that array.

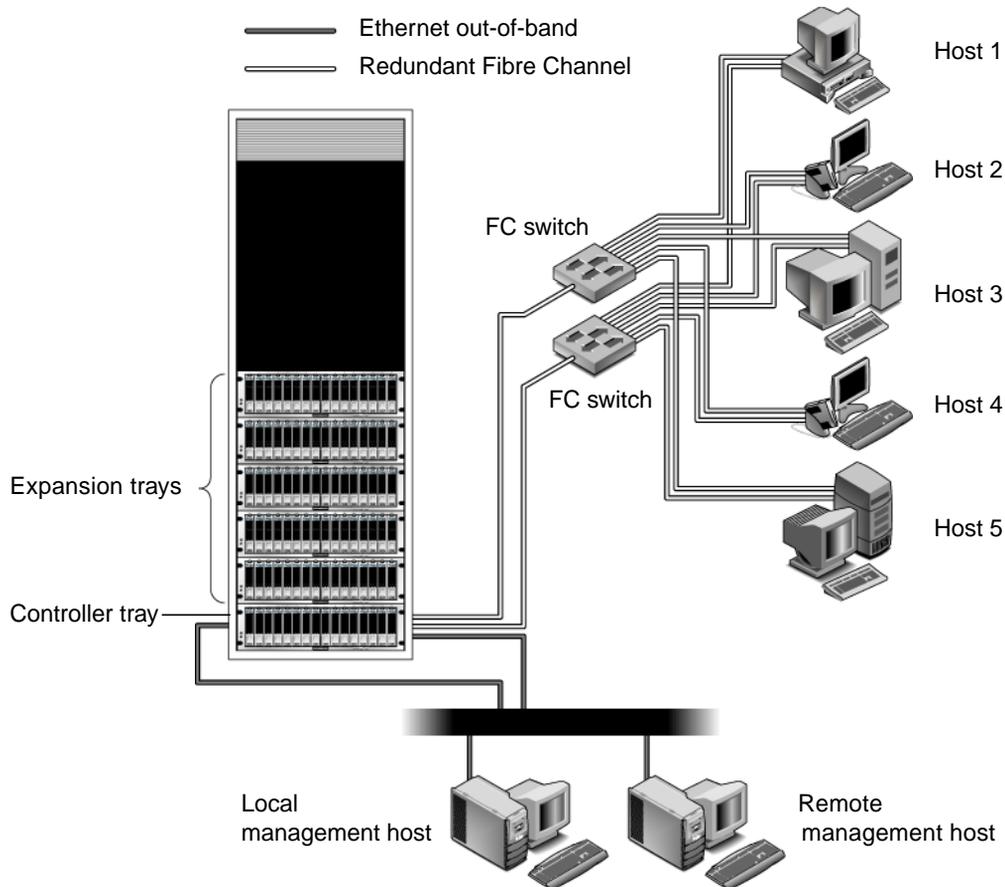


FIGURE 1-1 Sun StorageTek 6140 Array Product OverviewSun StorageTek 6140 Array Product Overview

About the Array Software

The array software is delivered on compact disk (CD) and consists of:

- "Management Software" on page 4
- "Remote CLI Client" on page 5
- "Monitoring and Diagnostic Software" on page 5
- "Data Host Software" on page 5

- “Optional Software” on page 6
- “Management-Related Service” on page 6

You specify the functionality you require, and the compact CD installs the necessary software.

Management Software

The web-based management software is the primary interface for configuring and managing the array. The management software consists of a suite of tools that you install on an external management host. The management host must be a Sun workstation running the Solaris 8, Solaris 9, or Solaris 10 Operating System (OS).

The management software enables the storage administrator to manage the array from any system with a web browser that is on the same network as the management host. For a list of supported browsers, see the release notes.

The management software, labeled Sun StorageTek Configuration Service in the graphical user interface (GUI), provisions and maintains the storage for data hosts.

- **Storage pools** are collections of volumes that share a profile defining the common configuration of the volumes.
- **Storage profiles** define the characteristics of a storage pool. You can choose one from the set of preconfigured profiles or create a new one.
- **Volumes** are divisions of a storage pool, consisting of virtual disks, representing the storage space that is used by the data hosts in the environment.
- **Virtual disks**, also called redundant array of independent disks (RAID) sets, are collections of locations in the memory of more than one physical disk. The storage array handles a virtual disk as if it were an actual disk.
- **Host groups** are a collection of hosts that will share access to the same volumes.
- **Snapshots** are copies of the data in a volume. They offer a high-availability alternative to backups because you do not need to take the array offline to create the snapshot, and you can store the snapshots in less space than the original data.
- **Volume copies** are copies of the complete contents of one data volume that are located on another data volume on the same array.
- **Replication sets** are the association between primary and secondary volumes. The secondary volume contains a complete copy of the data on the primary volume. The data replication software continuously replicates the data between volumes in a replication set.

Remote CLI Client

You can also manage and configure storage for the array using the remote command line interface (CLI) client. The CLI provides the same control and monitoring capability as the web browser, and it is also scriptable for running frequently performed tasks.

The remote CLI client is available for Solaris operating system (OS) and several other operating systems. See the release notes for a list of supported operating system platforms. For more information about the CLI commands, see the `sscs` man page.

Monitoring and Diagnostic Software

The Sun Storage Automated Diagnostic Environment is a monitoring and diagnostic tool for the array. You can configure the software to monitor on a 24-hour basis, collecting information that enhances the reliability, availability, and serviceability (RAS) of the array.

The Sun Storage Automated Diagnostic Environment can be accessed from a web browser or from the command line. It automates the transmission of alerts, which can be sent to an email address, to a pager, or to diagnostic software running on a management host on the network.

The monitoring software records alerts and notifications. You can monitor events and alerts by displaying the log file. You can also set up event notification so that notifications of actionable events are sent to your email address or pager.

Data Host Software

The array data host software controls the data path between the data host and the array. The data host software consists of the following tools:

- Sun StorEdge SAN Foundation Software for managing the data path I/O connections between data hosts and the array. This software includes drivers and utilities that enable Solaris data hosts to connect to, monitor, and transfer data in a Storage Area Network (SAN).
- Sun StorEdge Traffic Manager software which provides multipathing functionality and the ability to reliably communicate with the array's storage.

Data host software enables Solaris 8, Solaris 9, and Solaris 10 workstations and other operating system platforms to communicate with the array. For a list of supported operating system platforms, see the release notes.

Data host software for Solaris is distributed on the host software installation CD. You can also obtain the software for other operating systems from Sun's Download Center.

Optional Software

You can also purchase the following software for use on hosts with data paths or network connections to the array:

- Sun StorageTek Enterprise Storage Manager diagnostic software, for installation on a management host
- Sun StorageTek Enterprise Storage Manager agent, for installation on a Solaris or third-party data host that also has an out-of-band management network connection to the array
- VERITAS Volume Manager with Dynamic Multipathing (DMP), for installation on a Solaris or third-party data host (DMP can coexist with Sun StorEdge Traffic Manager multipathing software)
- Sun StorEdge SAM-FS software for remote archival and transparent backups
- Sun Cluster or Microsoft clustering software for clustered configurations

Management-Related Service

You can manage the array using the Hypertext Transfer Protocol Secure (HTTPS) web interfaces for management and monitoring of storage. HTTPS is primarily for web browsers and is protected by passwords and data encryption. The remote command-line interface (CLI) client on a management host is also used for management and monitoring of storage.

Hardware Components Overview

This section describes array hardware. It contains the following subsections:

- “About High-Availability Features” on page 7
- “About the Sun StorageTek 6140 Storage Array Hardware” on page 8
- “About the Sun StorageTek 6130 Storage Array Hardware” on page 13

About High-Availability Features

The Sun StorageTek arrays, including the 6130, 6140, and the 6540, are designed to provide high levels of reliability, availability, and serviceability (RAS). All components in the array’s data path are redundant so that the data path has no single point of failure. If one component fails, the array fails over to the alternate component automatically. The following components have been designed for high-availability:

- **Storage trays** – The redundant array of independent disks (RAID) controllers operate in master/alternate master mode.
- **Hot-spare drives** – Hot-spare drives may be allocated from unused drives and are always powered up and available but are not part of the array’s virtual disk. An array hot-spare is available as a spare to any virtual disk in any tray in the array configuration.

A configuration that has hot-spares in every storage tray supports high availability.

The array 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 array. The array generates an error event and records it in log files. Error events and log files are managed by the Sun Storage Automated Diagnostic Environment.

Refer to the release notes for information specific to that array.

About the Sun StorageTek 6140 Storage Array Hardware

This section describes the Sun StorageTek 6140 array hardware. It contains the following subsections:

- “Hardware Overview” on page 8
- “About the Sun StorageTek 6140 Controller Tray” on page 9
- “About the Sun StorageTek 6140 Expansion Tray” on page 12

Hardware Overview

The Sun StorageTek 6140 Array is a modular storage device that can contain up to 16 disk drives. A maximum of eight trays, one controller tray and seven expansion trays, can be installed in the Sun StorEdge Expansion cabinet, Sun Fire cabinet, and Sun Rack 900 cabinets.

For redundant power, each array controller has two power supplies, each with its own battery backup system. In a total power failure, each array has sufficient power from the batteries to shut down in an orderly fashion. In addition, each disk drive has array information that is mirrored across all drives for redundancy so that data can be recovered from any functional drive.

The array includes the following components:

- **Power supplies** – The two power supplies are housed in removable components. The power supplies convert incoming AC power to DC power. They are hot swappable and provide redundancy. It takes one power supply to power a fully loaded chassis. The power supplies input voltage range is 50 to 60 cycles per second.
- **Fan assembly** – Each array has two removable fan components, each containing two fans. The fans draw ambient air into the chassis through the front intake vents and push the air out the vents on the back of each fan. The fans provide redundant cooling, which means that if one of the fans in either component fails, the remaining fans will continue to provide sufficient cooling to operate the array. The fan assembly is hot-swappable from the rear of the array chassis.

About the Sun StorageTek 6140 Controller Tray

A controller tray contains two redundant array of independent disks (RAID) controllers, which operate independently and provide failover capability for the management and data paths. The controller tray is configured for Fibre Channel (FC) disk drives and provides RAID functionality, caching, and disk storage.

The RAID controller has 1 gigabyte of memory for processor memory and data cache. The processor memory is used to store application data, while the data cache is a buffer used to store data temporarily during data read and write functions.

Each controller contains a sealed, rechargeable 4-volt, lead acid battery. This battery provides backup power to cache memory for up to three days in the event of a power loss. You must replace the battery every two years. The battery is customer replaceable and can be accessed from the rear of the controller on the back of the array tray.

The Sun StorageTek 6140 array controller tray has the components listed in TABLE 1-1.

TABLE 1-1 Sun StorageTek 6140 Array Controller Tray

Description	Quantity
FC RAID controllers	2
FC disk drives	5 to 16 4-GB or 2-GB drives per tray
Ethernet ports for management host connections	4 (2 per controller)
4/2/1 Gbps FC host ports with SFPs	8 (4 per controller)
4/2 Gbps FC expansion ports	4 (2 per controller)
Power supplies	2
Fan assemblies	2

FIGURE 1-2 shows the LEDs and components at the front of the controller tray.

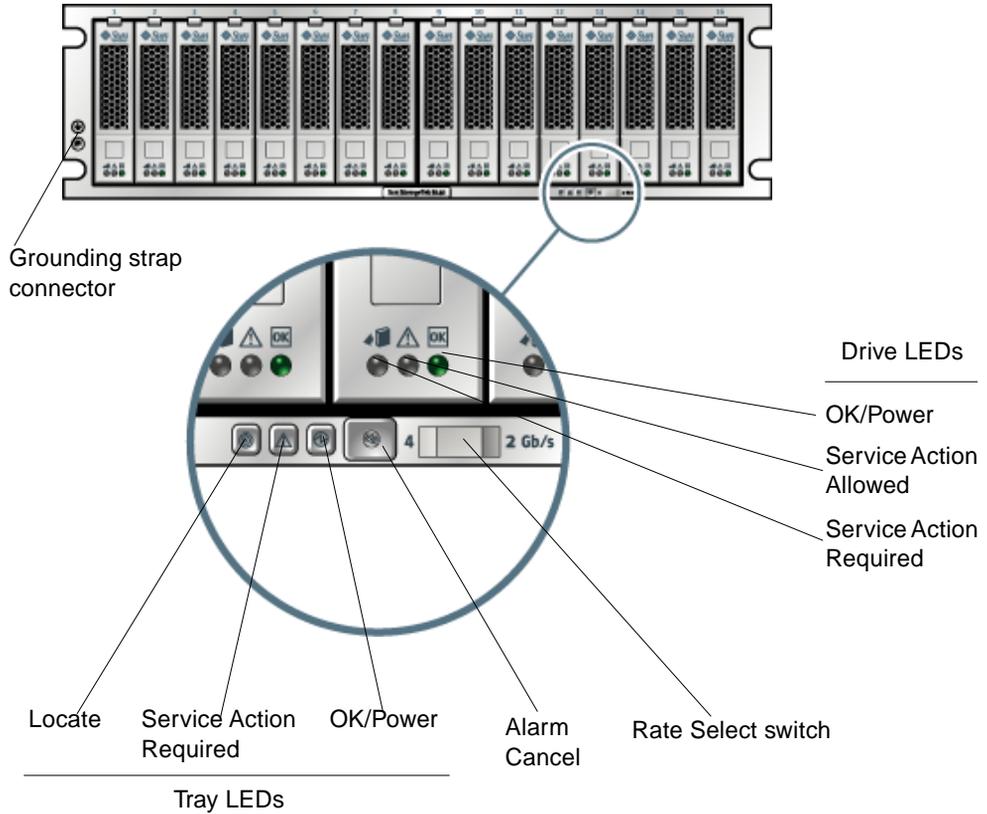


FIGURE 1-2 Sun StorageTek 6140 Array Controller Tray (Front View)

FIGURE 1-3 shows the ports and switches at the back of the Sun StorageTek 6140 array controller tray.

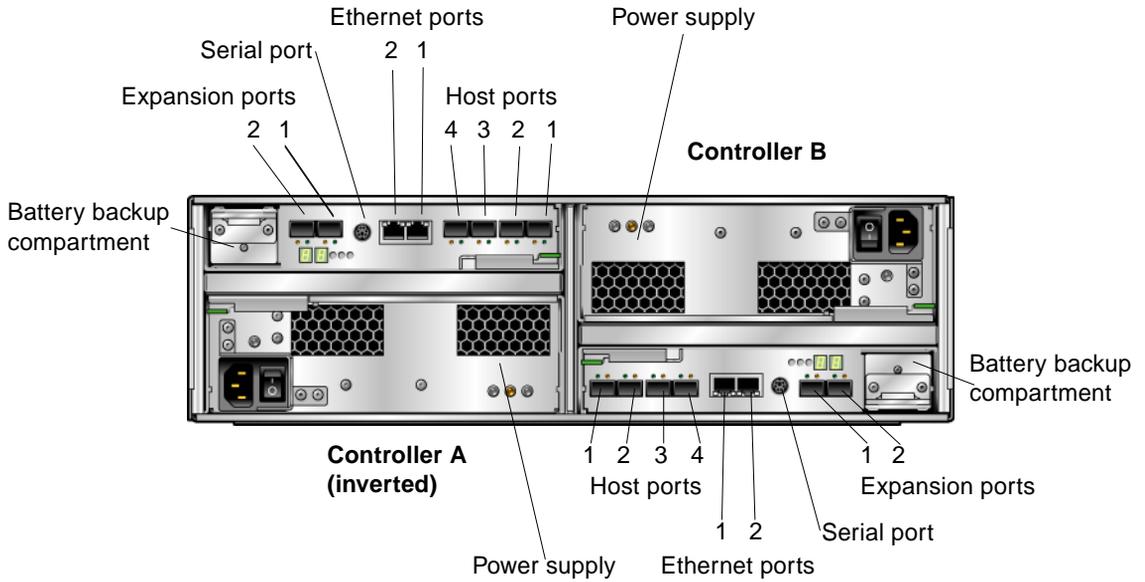


FIGURE 1-3 Sun StorageTek 6140 Array Controller Tray Ports and Switches (Back View)

TABLE 1-2 describes the ports and switches on the back of the Sun StorageTek 6140 array controller tray.

TABLE 1-2 Sun StorageTek 6140 Array Controller Tray Ports and Switches (Back)

Ports/Switches	Description
Host ports (Ch1 - Ch4)	Four 4-, 2-, or 1-Gbit/second FC Small Form-factor Plug-in (SFP) ports. Ch4 host port is reserved for remote replication requirements.
Ethernet ports (1 and 2)	RJ-45 Ethernet ports used for out-of-band management of the RAID controller. An internal Ethernet device provides standard 10 Mbits/second and 100 Mbits/second full-duplex connectivity.
Expansion ports (P1 and P2)	2-Gbit FC ports used to connect to the drive channel device and expansion trays.
Serial port	A port that allows terminal access to display or configuration of the IP address for the tray, and for recovery of a lost password for the tray.

About the Sun StorageTek 6140 Expansion Tray

The expansion tray provides from 5 to 16 additional FC or Serial Advanced Technology Attachment (SATA) drives. An expansion tray is cabled directly to a controller tray and cannot operate independently.

TABLE 1-3 describes the Sun StorageTek 6140 array expansion tray configuration.

TABLE 1-3 Sun StorageTek 6140 Array Expansion Tray

Description		Quantity
FC or SATA disk drives	FC hard disk drives 73G10K, 73G15K, 146G10K	5 to 16 4- or 2-Gbit/second drives
	SATA hard disk drives :400G7.2	5 to 16, 4- or 2-Gbit/second drives
Drive expansion ports		4
Power supplies		2
Fan assemblies		2

FIGURE 1-4 shows the ports and components on the back of the Sun StorageTek 6140 array expansion tray.

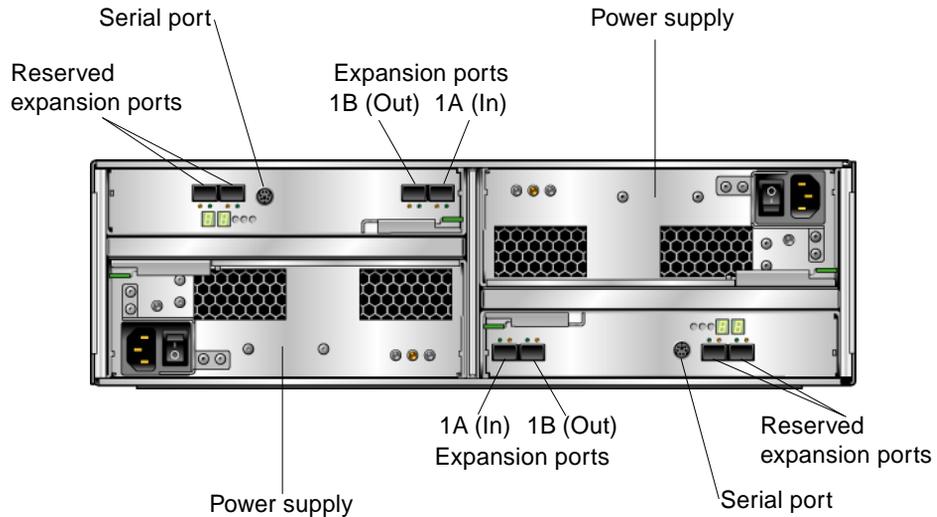


FIGURE 1-4 Sun StorageTek 6140 Array Expansion Tray Ports and Components (Back)

TABLE 1-4 describes the ports and components on the back of the Sun StorageTek 6140 array expansion tray.

TABLE 1-4 Sun StorageTek 6140 Array Expansion Tray Ports and Components (Back)

Ports/Switches/LEDs	Description
Expansion ports 2A (In), 2B (Out)	Two 4- or 2-Gbit FC ports used to connect to an array controller and/or additional expansion trays.
Serial port	Allows terminal access to display or configure the IP address for the tray and to recover a lost password for the tray.

About the Sun StorageTek 6130 Storage Array Hardware

This section describes the Sun StorageTek 6130 hardware. It contains the following subsections:

- “Hardware Overview” on page 13
- “About the Sun StorageTek 6130 Controller Tray” on page 14
- “About the Sun StorageTek 6130 Expansion Tray” on page 15

Hardware Overview

The Sun StorageTek 6130 array is a modular storage device that can contain up to 14 disk drives. A maximum of eight arrays, one controller tray and seven expansion trays, can be installed in the Sun StorEdge Expansion cabinet, Sun Fire cabinet, and Sun Rack 900/1000 cabinets.

For redundant power, each array has two power supplies, each with its own battery backup system. In a total power failure, each array has sufficient power from the batteries to shut down in an orderly fashion. In addition, each disk drive has array information that is mirrored across all drives for redundancy so that data can be recovered from any functional drive.

The array includes the following components:

- **Power supplies** – The two power supplies are housed in removable components. The power supplies convert incoming AC power to DC power. They are hot swappable and provide redundancy. It takes one power supply to power a fully loaded chassis.
- **Fan assembly** – Each array has two removable fan components, each containing two fans. The fans draw ambient air into the chassis through the front intake vents and push the air out the vents on the back of each fan. The fans provide redundant cooling, which means that if one of the fans in either component fails, the remaining fans will continue to provide sufficient cooling to operate the array. The fan assembly is hot-swappable from the rear of the array chassis.

About the Sun StorageTek 6130 Controller Tray

A controller tray contains two redundant array of independent disks (RAID) controllers, which operate independently and provide failover capability for the management path. The controller tray is configured for Fibre Channel (FC) disk drives and provides RAID functionality, caching, and disk storage.

The RAID controller has 1 gigabyte of memory for processor memory and data cache. The processor memory is used to store application data, while the data cache is a buffer used to store data temporarily during data read and write functions.

Each controller contains a sealed, rechargeable 4-volt, lead acid battery. This battery provides backup power to cache memory for up to three days in the event of a power loss. You must replace the battery every two years.

The Sun StorageTek 6130 array controller tray has the components listed in TABLE 1-5.

TABLE 1-5 Sun StorageTek 6130 Array Controller Tray

Description	Quantity
FC RAID controllers	2
FC 3.5-inch drives	5 - 14 per tray
FC hard disk drives 73G10K, 73G15K, 146G10K, 146G15K, 300G10K with Small Form-factor Plug-ins (SFPs) 10K	
Ethernet ports for management host connections	2 (1 per controller)
2-gigabits per second (Gbps) FC host ports with SFPs	4 (2 per controller)
2-Gbps FC expansion ports	2 (1 per controller)

TABLE 1-5 Sun StorageTek 6130 Array Controller Tray (Continued)

Description	Quantity
Power supplies	2
Fan assemblies	2

Light emitting diodes (LEDs) and switches on the front of the controller tray are available to monitor the state of the controller. The ports, switches, and LEDs at the back of the controller tray also enable you to monitor state of the controller and are especially useful during installation.

FIGURE 1-5 shows the LEDs on the front view of the Sun StorageTek 6130 controller tray.

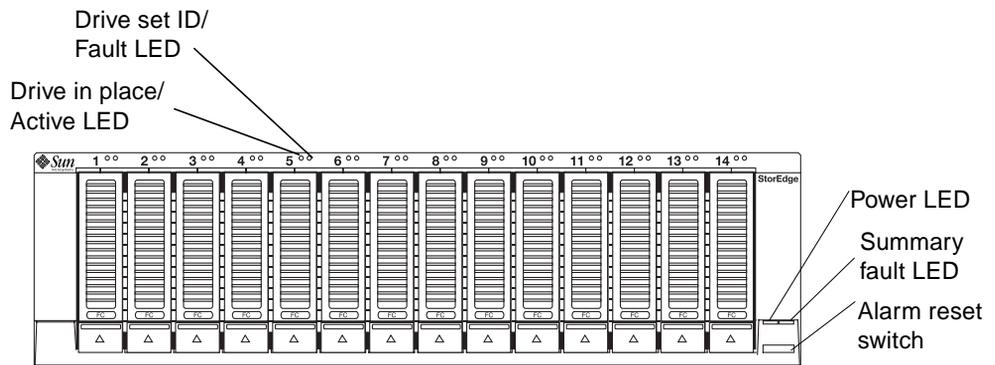


FIGURE 1-5 LEDs on the Front of the Sun StorageTek 6130 Array Controller Tray

About the Sun StorageTek 6130 Expansion Tray

The expansion tray provides for up to 14 additional Fibre Channel (FC) or Serial Advanced Technology Attachment (SATA) drives. An expansion tray is cabled directly to a controller tray and cannot operate independently.

You cannot mix FC and SATA disk drives within the same expansion tray. All drives within a single expansion tray must be the same type, either all FC drives or all SATA drives.

The Sun StorageTek 6130 array controller tray has the components listed in TABLE 1-6.

TABLE 1-6 Sun StorageTek 6130 Array Expansion Tray

Description	Quantity
FC or SATA 3.5-inch drives	5 - 14
FC hard disk drives: 73G10K, 73G15K, 146G10K, 146G15K, 300G10K	
SATA model: 400G7.2K	
I/O tray with In and Out expansion ports	2
Power supplies	2
Fan assemblies	2

FIGURE 1-6 shows the front and back views of the Sun StorageTek 6130 array expansion tray.

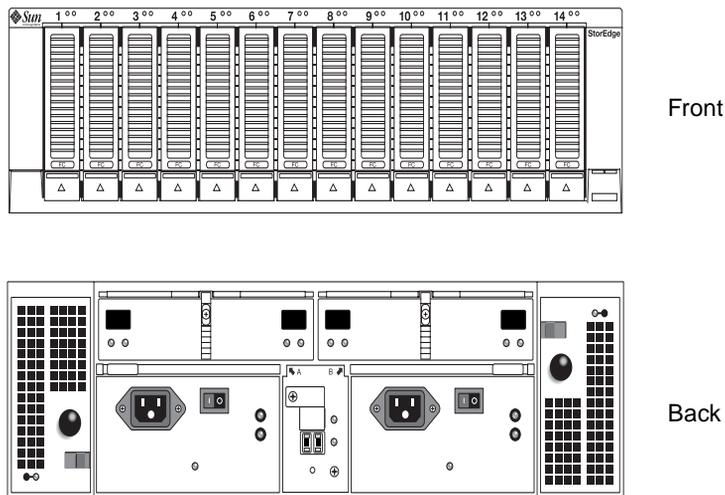


FIGURE 1-6 Sun StorageTek 6130 Array Expansion Tray - Front and Back Views

Using the Browser Interface

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

- “About the Browser Interface” on page 18
- “About the Page Banner” on page 18
- “About the Navigation Tree” on page 19
- “About the Page Content Area” on page 20
- “Controlling the Display of Table Information” on page 21
- “About the Status Icons” on page 22
- “Using Forms” on page 23
- “Searching for System Elements” on page 24
- “Using Help” on page 25
- “Logging Out of the Management Software” on page 26

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 navigation tree 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, tree object, 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 Banner

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

TABLE 2-1 describes the contents of the banner.

TABLE 2-1 Contents of the Banner

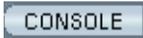
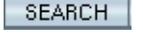
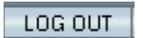
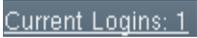
Button	Description
	Returns you to the Java Web Console page, where you can navigate between the configuration software and the diagnostic software.
	Displays the software version and copyright information.
	Refreshes the current page.
	Enables you to quickly locate logical and physical elements defined 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 24.
	Logs you out of the Java Web Console and the current application.

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

Button	Description
	Opens the online help in a separate window.
System Information and Status	
	Displays the name of the user who is currently logged in to the system.
	Displays the name of the system.
	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.
	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.
	Displays the current number of each type of alarm. There are four alarm types:  Down,  Critical,  Major, and  Minor. To get more information about the alarms, click the Current Alarms link. When you click Current Alarms from a configuration service window, the diagnostic environment is launched in a separate window and the Alarm Summary page is displayed.

About the Navigation Tree

The navigation tree is displayed in the left-hand pane of the interface. You use the navigation tree to move among folders and pages within an application.

The top of the navigation pane displays the Change Array link. Click this link to return to the Array Summary page, where you can select a different array to manage.

TABLE 2-2 describes the top-level objects in the Sun StorageTek Configuration Service navigation tree.

TABLE 2-2 Sun StorageTek Configuration Service Tree - Top-Level Components

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.
Mappings	Enables you to view system-wide mappings.
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 a link in the page to perform a task or to move among pages. You can also move among pages by clicking an object in the navigation tree.

Controlling the Display of Table Information

Tables display data in a tabular format. TABLE 2-3 describes the objects you can use to control the display of data on a page.

TABLE 2-3 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>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>
	<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>

TABLE 2-3 Table Objects (*Continued*)

Control/Indicator	Description
	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). Click this icon to change the sort order of the column to ascending. A closed icon indicates the column by which the table is currently sorted.
	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.
	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.
	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.

About the Status Icons

Icons are displayed to draw your attention to an object's status. TABLE 2-4 describes these status icons.

TABLE 2-4 Status Icons

Control/Indicator	Description
	Identifies a critical error. Immediate attention to the failed object is strongly recommended.
	Identifies a a minor error. The object is not working within normal operational parameters.

TABLE 2-4 Status Icons (*Continued*)

Control/Indicator	Description
	Identifies an unknown condition. A report on the status cannot be supplied at this time.

Using Forms

Forms have menus, buttons, links, and text fields that allow you to select available options and enter information on a page. TABLE 2-5 describes these elements.

TABLE 2-5 Form Controls

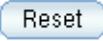
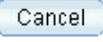
Control/Indicator	Description
*	Indicates that you must enter information in this field.
	Lists options from which you can make a selection.
	Displays the part of the form that is indicated by the text next to this icon.
	Returns you to the top of the form.
	Saves the selections and entries that you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

TABLE 2-5 Form Controls (Continued)

Control/Indicator	Description
	Causes the current settings to take effect.

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 World Wide Name (WWN).

To use the search feature:

1. Click Sun StorageTek 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, hosts, host groups, volumes, replication sets, snapshots, 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 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 describes the help tabs.

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.

TABLE 2-7 describes meanings of the Help window icons.

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 performing connection and administration tasks. It contains the following subsections:

- “Connection Tasks” on page 28
- “Administration Tasks” on page 31

Connection Tasks

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

- “Managing From the Site LAN” on page 28
- “Managing an Array Isolated From the LAN” on page 29
- “Installing the Remote CLI Client” on page 29
- “About Host Bus Adapters” on page 30
- “About Multipathing” on page 30

Managing From the Site LAN

You can manage the storage in the array with a browser from any host that has a network connection to the management host.

An Ethernet cable connects your management host in the site local area network (LAN) to the array. For more information on cable connections and setting Internet Protocol (IP) address, see the Getting Started Guide for your array.

You have the following options in how you connect the array to the site LAN:

- Dynamic IP address. The array can get its IP addresses from your site’s Dynamic Host Configuration Protocol (DHCP) server each time it logs in to the LAN.
- Static IP address. You can set a static IP addresses for the master and alternate master controllers.
- Default IP address. You can use default IP addresses for the master and alternate master controllers.
- The level of access that the array allows from the network.

You can also manage the array and monitor, diagnose, and fix problems manually with `sscs` commands or by means of scripts using a remote command-line interface (CLI) client. Remote CLI clients are available for Solaris, Microsoft Windows, LINUX, IBM AIX, and HPUX operating systems.

Managing an Array Isolated From the LAN

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

- Set up and configure the array using a management host.
- When you have finished configuring the array, disconnect the management host.

Whenever the array 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 array and the external LAN.

Installing the Remote CLI Client

If you need to configure the array with the `sscs` commands from a host other than the management host, the host installation software compact disk (CD) contains a remote command-line interface (CLI) that enables you to do all 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, XP

To install the remote CLI on a Solaris host:

1. Insert the Host Installation Software CD.
2. Run the `install` 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 Getting Started Guide for your array.

About Host Bus Adapters

A data host is any host that uses the array for storage. When a data host is connected to the array by a host bus adapter (HBA), the HBA is an initiator. The HBA is connected by a cable to a Fibre Channel (FC) port on the array.

After connecting a data host directly to the array with one or two HBAs, use the `luxadm(1M)` command to verify the firmware level of HBAs. If the firmware is not at the correct revision, use the host installation software compact disk (CD) to install the Sun StorageTek SAN Foundation software. You can then configure the initiator and set up hosts and host groups.

About Multipathing

With multipathing, also called multipath failover, an array 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 array, 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, also called MPxIO, is Sun's multipathing solution for Fibre Channel connected storage devices. This software is provided as part of the Sun StorEdge SAN Foundation software. See the SAN Foundation software documentation for information on downloading, installing, and configuring this software.
- VERITAS Volume Manager with Dynamic Multipathing (DMP)

Administration Tasks

This section describes performing administration tasks for the array. It includes the following subsections:

- “Accessing, Adding, and Deleting Storage Arrays” on page 31
- “Specifying General Settings” on page 34
- “Setting the Array IP Address” on page 38
- “Managing User Accounts” on page 39
- “Managing Licenses” on page 42
- “Displaying Event Information” on page 44
- “Monitoring Performance” on page 45
- “Administering Jobs” on page 45
- “Viewing the Activity Log” on page 46

Accessing, Adding, and Deleting Storage Arrays

When you log in to the management software, the Array Summary page lists the available arrays.

To manage an existing array, click it. This gives you access to logical and physical components associated with the array, including replication sets, host groups, hosts, initiators, storage pools, volumes, virtual disks, trays, and disks.

You can add additional arrays by auto-discovery or manually by individual registration.

You can also remove existing arrays.

Displaying Array Information

To display available arrays:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click an array name to view additional information about that array.

The navigation pane and the Volume Summary page for the selected array are displayed.

Planning to Register an Array

Using the Array Registration wizard, you can have the management software either auto-discover one or more arrays that are connected to the network and are not already registered, or you can choose to manually register an array.

The auto-discover process sends out a broadcast message across the local network to identify any unregistered arrays. The discovery process displays the percentage of completion while the array management software polls devices in the network to determine whether any new arrays are available. When complete, a list of discovered arrays is displayed. You can then select one or more arrays to register from the list.

Manual registration enables you to register an array by identifying the IP address of its controller. This option is typically used only to add a storage array that is outside of the local network.

The Array Registration wizard displays firmware information for each array and, for each Sun StorageTek 6140 and 6130 array, lists any action recommended to bring the array up to the current firmware baseline level. You can choose to perform the recommended firmware upgrade action now, or you can modify the array firmware later by selecting the array and clicking the Upgrade Firmware button on either the Array Summary page or the Administration > General page.

For information about how to upgrade the firmware for the Sun StorageTek 6540 array, see the release notes.

Registering an Array

To register an array:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click Register.

The management software launches the Register Array wizard.

3. Follow the instructions in the wizard.

Unregistering an Array

To unregister an array:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Select the check box to the left of the array you want to remove from the list of registered arrays.
This enables the Remove button.
3. Click Remove.

Upgrading Array Firmware

For optimal performance, Sun Microsystems recommends that the firmware on all arrays be at the level of the current firmware baseline. For the Sun StorageTek 6140 and 6130 arrays, you can upgrade array firmware during array registration. To upgrade the firmware for the Sun StorageTek 6540 array, see the instructions in the *Sun StorageTek 6540 Release Notes*.

To upgrade the array firmware:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Select the check box to the left of one of the displayed arrays.
This enables the the Upgrade Firmware button.
3. Click Upgrade Firmware.
The management software launches the Upgrade Firmware wizard.
4. Follow the instructions in the wizard.

Note: You can upgrade the firmware during array registration and you can also launch the Upgrade firmware wizard from the General Setup page.

Monitoring Array Health

To monitor the health of the array:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Note the value in the Health field, which indicates the state of the array. Possible values include:

- **OK** - Indicates that every component of the storage array is in the desired working condition.
- **Degraded** - Indicates that, although the management host can communicate with the storage array, a problem on the array requires intervention. For example, the array may have volumes that are not on the array's preferred I/O controller path. To correct this problem, click the Redistribute Volumes button. If the array status remains in the Degraded state, access the Sun Storage Automated Diagnostic Environment for further information on troubleshooting and correcting this problem.

Typically, multipath drivers move volumes from their preferred owner controller when a problem occurs along the data path between the host and the storage array. Redistributing the volumes causes the volumes to revert back to their preferred controllers.

Redistributing the volumes while an application is using the affected volumes causes I/O errors unless a multipath driver is installed on the data host. Therefore, before you redistribute volumes, verify either that the volumes are not in use or that there is a multipath driver installed on all hosts using the affected volumes.

- **Error** - Indicates that the management host cannot communicate with the controllers in the storage array over its network management connection. Access the Sun Storage Automated Diagnostic Environment for information on troubleshooting and correcting this problem.

Specifying General Settings

You use the General Setup page to manage passwords, to view and specify array details, to enable disk scrubbing on the array, and to set the time on the array. Any user can view the information on this page, but you must be logged in as a user with storage role to change the settings.

About the Array Password

An array password provides access to an array and is required for performance of notification operations. When you set the password, the management software stores an encrypted copy of the password in its array registration database. Thereafter, the management software can perform modification operations on the array without a password challenge.

You can change the array password at any time. Changing the array password causes the management software to automatically update the password stored in its array registration database.

Multiple management hosts can access a single array. Each management host has its own instance of management software, each of which has its own array registration database. For the management software to perform modification operations on an array, the password stored in the array registration database for that instance of the management software must match the password set on the array. When you change the array password on one management host, only the array registration database used by that management host is updated with the changed password. Before another management host can perform modification operations on that array, the array registration database for that management host must also be updated with the new password.

You may also need to update the password stored in the array registration database if the array was registered without a password or with an incorrectly typed password.

If the password stored in the array registration database does not match the array password, the following error message is displayed when you attempt a modification operation on the array: "The operation cannot complete because you did not provide a valid password."

Changing the Array Password

To change the array password or update the password stored in the array registration database:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to change the password.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > General Settings.
The General Setup page is displayed.
4. Click Manage Passwords.
The Manage Passwords page is displayed.
5. Select one of the following:
 - To change the array password, select Change Array Password. Changing the password automatically updates the password stored in the array registration database.

- To manually synchronize the password stored in the array registration database with the password set on the array, select Update Array Password In Array Registration Database. Do this if the array password was previously changed from another management host or, if the array was registered without a password or with an incorrectly typed password.
6. In the Old Password field (available only if you are changing the array password), enter the current password.
 7. In the New Password field, enter the new password as an alphanumeric string of up to eight characters.
 8. Enter the same new password in the Verify New Password field.
 9. Click OK to apply your changes.

Setting Array Details

You can view details about the array, and you change some, including the array name, the number of hot-spares, the default host type, the cache block size, the minimum and maximum cache allocation percentages, disk scrubbing, and failover alert settings.

To set array details:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to set the IP address.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > General Settings.
The General Setup page is displayed.
4. Go to the Details section.
5. Specify the new settings you want to change.
6. Click OK to apply your changes.

Enabling Disk Scrubbing

Disk scrubbing is a background process performed by the array controllers to provide error detection on the drive media. Disk scrubbing detects errors and reports them to the event log.

Before disk scrubbing can run, you must enable it on the array. Disk scrubbing then runs on all volumes on the array. You can disable disk scrubbing on any volume that you do not want to have scrubbed. Later, you can re-enable disk scrubbing for any volume on which you disabled it.

The advantage of disk scrubbing is that the process can find media errors before they disrupt normal drive reads and writes. Disk scrubbing scans all volume data to verify that it can be accessed. If you enable a redundancy check, it also scans the volume redundancy data.

Enabling Disk Scrubbing on an Array

To enable disk scrubbing on an array:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to enable disk scrubbing.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Administration > General Settings.
The General Setup page is displayed.
4. Click the check box next to Disk Scrubbing Enabled and specify the number of days for a disk scrubbing cycle.
5. Click OK.

Disabling and Re-enabling Disk Scrubbing on a Volume

To disable or re-enable disk scrubbing for an individual volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to disable or re-enable disk scrubbing for a specific volume.
The navigation pane and the Volume Summary page are displayed.
3. Click the volume on which you want to disable or re-enable disk scrubbing.
The Volume Details page is displayed.

4. Do one of the following:
 - To disable disk scrubbing, select False in the Disk Scrubbing Enabled field.
 - To re-enable disk scrubbing, select True in the Disk Scrubbing Enabled field.
 - To re-enable disk scrubbing so that it also scans the volume redundancy data, select True in the Disk Scrubbing With Redundancy field.
5. Click OK.

Setting the Time

If the array does not use your network's network time protocol server, you must set the array's clock manually.

To set the array time:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to set the time.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > General Settings.
The General Setup page is displayed.
4. Go to the System Time section.
 - To synchronize the array time with the server, click Synchronize With Server.
 - To set the time manually:
Set the hour and minute, using a 24-hour clock.
Set the month, day, and year.
5. Click OK to apply your changes.

Setting the Array IP Address

You must specify the method by which an Internet Protocol (IP) address is supplied for the array. If you choose the Dynamic Host Control protocol (DHCP), the network provides an Internet Protocol (IP) address for the array each time the array is

powered on and logs in to the network. Alternatively, you can choose a static IP address, which the array will use every time it is powered on and logs in to the network.

Note: The web browser you use to manage the array relies on the array's IP address. If the address changes because either you changed it manually or the system was assigned a new one, the browser loses its connection to the array. You must reconnect to the array to continue monitoring and managing the array.

To set the IP address:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to set the IP address.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Physical Storage > Controllers.
The Controller Summary page is displayed.
4. In the Ethernet Port 1 field, select Enable DHCP/BOOTP or Specify Network Configuration. If you select Specify Network Configuration, you must enter an IP address, gateway address, and netmask of the controller using Ethernet port 1.
5. (6140 array only) In the Ethernet Port 2 field, select Enable DHCP/BOOTP or Specify Network Configuration. If you select Specify Network Configuration, you must enter an IP address and netmask of the controller using Ethernet port 2.
6. Click OK.

Managing User Accounts

This section describes managing user accounts. It contains the following subsections:

- "About User Accounts" on page 40
- "Displaying User Information" on page 40
- "Adding a New User" on page 41
- "Removing a User" on page 41

About User Accounts

The management software and the Sun Storage Automated Diagnostic Environment software installed on the management host share user roles that define the privileges available to the user. TABLE 3-1 describes the user roles and their privileges.

TABLE 3-1 User Roles and Privileges

Role	Description of Role
storage	Users assigned the storage role can view and modify all attributes.
guest	Users assigned the guest role can view all attributes but not modify any of them.

After installing the array software on a server and logging in to the configuration services software using root, you can assign one of the roles to valid Solaris user accounts that have access to the management host. The users can then log in to the configuration services software using their Solaris user names and passwords. For information about creating Solaris user accounts, refer to the Solaris system administration documentation.

The role assigned to a user determines the degree of access that a user has to the array and its attributes. All users assigned the guest role can view information. To modify array attributes, a user must have storage privileges. Only users assigned the storage role can add users that have either the guest or storage account role.

If multiple users are logged in to the array and making changes as the storage administrator, there is a risk of one user's changes overwriting another user's previous changes. Therefore, storage administrators should develop procedures about who can make changes and when and about how to notify others.

Displaying User Information

To display user information:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display user information.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > User Management.
The User Summary page is displayed.

Adding a New User

To add a new user:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to add a new user.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > User Management.
The User Summary page is displayed.
4. Click Add to add a new user.
The Add New User page is displayed.
5. Enter the name of a valid Solaris account.
Solaris user names should begin with a lowercase letter and can consist of 6 to 8 alphanumeric characters, underscores (_), and periods (.). For information about creating Solaris user accounts, refer to the Solaris system administration documentation.
6. Select a role for the Solaris account: `storage` or `guest`.
7. Click OK.
The User Summary page lists the newly assigned user and role.

Removing a User

To remove a user:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to remove a user.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > User Management.
The User Summary page is displayed.
4. Click the check box to the left of the name of the user you want to remove.
This enables the Remove button.

5. Click Remove.

Note: For information about removing user accounts from the server or NIS, consult the Solaris system administration documentation.

Managing Licenses

This section describes managing licenses. It contains the following subtopics:

- “About Licensed Features” on page 42
- “Displaying License Information” on page 42
- “Adding a License” on page 43
- “Disabling a License” on page 43
- “Re-enabling a License” on page 44

About Licensed Features

Before you can use premium features, you must obtain and register licenses for each premium feature that you plan to use. Premium features include:

- Data Replication
- Volume Copy
- Snapshots
- Storage Domains

Displaying License Information

License certificates are issued when you purchase premium services and contain instructions for obtaining license information from the Sun Licensing Center.

To display license information:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display license information.
The navigation pane and the Volume Summary page for that array are displayed.

3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click a feature for detailed information on licenses for that feature.
The Licenseable Feature Details page for the selected feature is displayed.

Adding a License

License certificates are issued when you purchase premium services and contain instructions for obtaining license information from the Sun Licensing Center.

To add a license:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to add a new license.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Add License.
The Add License page is displayed.
5. Select the type of license you want to add, and specify the version number and key digest supplied by Sun.
6. Click OK.

Disabling a License

To disable a license:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to disable a license.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.

4. Click the check box to the left of the license that you want to disable, and click Disable.

Re-enabling a License

To re-enable a license, contact the Sun License Center (<http://www.sun.com/licensing>). Be prepared to provide the following information:

- Name of the product to be licensed
- Feature serial number, available from the license certificate
- Controller tray serial number, located on the back of the controller tray and on the Licensable Feature Summary page

Displaying Event Information

View events and configure notifications by opening the Notification Management page to access the Storage Automated Diagnostic Environment software.

To view event information and configure event notifications:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see event information.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Notification.
The Notification Management page is displayed.
4. Do one of the following:
 - To configure event notification, click Configure Notifications.
Note: You are required to enter at least one email address in the notification settings for events.
 - To view events, click Show Alarms.
The Storage Automated Diagnostic Environment interface is displayed in a separate browser window.

Monitoring Performance

To monitor array performance:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see performance statistics.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Performance Monitoring.
The Performance Monitoring page is displayed.
4. To turn performance monitoring on, select the Performance Monitoring Enabled check box and specify the polling interval you want.
5. To view current statistics, go to the Performance Statistics section of the page.

Administering Jobs

This section describes jobs. It contains the following subsections:

- “About Jobs” on page 45
- “Displaying Job Information” on page 46
- “Canceling Jobs” on page 46

About Jobs

When you request an operation on one object, the management software processes that operation immediately. For example, if you select one volume to delete, the volume is deleted immediately. However, because operations on more than one object can affect performance, when you request an operation on several objects, the system creates a job that completes the operation while you make other selections. You can follow the progress of a job from the Job Summary page.

Displaying Job Information

While a job is in progress, it is reported on the Job Summary page.

To display job information:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see job information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Jobs.
The Job Summary page lists the jobs that have been processed and their current status.
4. For more information about a job, click its identifier (ID).
The Job Details page for the selected job is displayed.

Canceling Jobs

You can cancel only volume copy jobs.

To cancel a job:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to cancel a job.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Jobs.
The Job Summary page is displayed.
4. Select the job you want to cancel, and click Cancel Job.

Viewing the Activity Log

The activity log lists user-initiated actions performed on the array, in chronological order. These actions may have been initiated through either the Sun StorageTek Configuration Service interface or the command-line interface (CLI).

To view the activity log:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display the activity log.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Administration > Activity Log.
The activity log is displayed.

Performing Configuration Tasks

This chapter describes configuration tasks. It contains the following subsections:

- “Initial State of the Array” on page 50
- “Configuring Storage Volumes” on page 53
- “Configuring Volume Copies” on page 62
- “Configuring Volume Snapshots” on page 70
- “Configuring Data Replication” on page 83
- “Configuring Virtual Disks” on page 106
- “Configuring Storage Pools” on page 110
- “Configuring Storage Profiles” on page 113
- “Configuring Initiators” on page 118
- “Configuring Host Groups and Hosts” on page 121
- “Managing Trays and Disk Drives” on page 132
- “Configuring Storage Domains” on page 139

Initial State of the Array

This section describes the initial state of the array and factors to consider before you configure the array. It contains the following subsections:

- “About the Initial State of the Array” on page 50
 - “About Provisioning Storage” on page 50
 - “About Physical Storage Elements” on page 51
 - “About Logical Storage Elements” on page 52
-

About the Initial State of the Array

After you complete the basic configuration tasks described in the Getting Started Guide for your array, all of the hardware and software is installed and at least one array is registered and named, and array passwords are set.

In addition, the following tasks will have been completed:

- The system time has been correctly set.
 - At least one new user has been added and assigned the role of “storage.”
 - An initiator has been created.
 - At least one host has been created and mapped to an initiator.
 - At least one host group has been created.
 - A storage pool has been created.
 - A volume has been created and mapped to a host or host group.
-

About Provisioning Storage

In a simple storage configuration, all data hosts could share all available storage in one storage pool, and any host mapped to an initiator would have access to any storage in the pool. Your organization’s needs determine whether you want a more complex storage configuration. For example, you might provision the storage for your organization by creating host groups and pools of virtual storage.

The array has a number of physical and logical storage elements that you can use to provision your storage:

- Physical storage elements: initiators, hosts, host groups, trays, and disks

- Logical storage elements: volumes, virtual disks, and pools

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

- **Security** – By creating host groups, you segregate initiators. For example, the hosts that handle financial data store their data in a different host group from the host group consisting of hosts that handle research data.
- **Input/output (I/O)** – Some storage profiles specify 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 array management software includes a set of profiles to meet various needs. You can also create custom profiles.

About Physical Storage Elements

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

- Trays hold the disk drives and support their operation. Each tray holds up to 14 disk drives.
- Disk drives are nonvolatile, randomly addressable, rewriteable data storage devices.
- Initiators are Fibre Channel (FC) ports on host bus adapters (HBAs) that allow hosts to gain access to a storage array.
- Hosts, or data hosts, are servers that can store data on a storage array. Data hosts are mapped to initiators.
- A host group is a collection of one or more hosts on one storage array that share access to the same volumes.

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 volumes that share a profile. The profiles defines the common configuration of the volumes.
- Virtual disks, also called redundant array of independent disks (RAID) sets, are a collection of locations in the memory of more than one physical disk. The storage array handles a virtual disk as if it were an actual disk. You create the virtual disks during volume creation.
- Volumes are divisions of a pool, consisting of virtual disks, and are accessed by hosts and host groups.
- Snapshots are copies of the data in a volume at a specific moment. Snapshots can be made without interruption of the normal operation of the system.

Configuring Storage Volumes

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

- “About Volumes” on page 53
 - “Planning Volumes” on page 54
 - “Managing Volumes” on page 56
-

About Volumes

You manage the array’s physical disks as a pool of storage space for creating volumes. Volumes are “containers” into which applications, databases, and file systems can put data. Volumes are created from the virtual disks, based on the characteristics of the storage pool associated with the virtual disks. Based on your specifications, the array automatically allocates storage a virtual disk that can satisfy your volume configuration requirements.

There are several different types of volumes:

- **Standard volume** – A standard volume is a logical structure created on a storage array for data storage. When you create a volume, initially it is a standard volume. Standard volumes are the typical volumes that users will access from data hosts.
- **Source volume** – A standard volume becomes a source volume when it participates in a volume copy operation as the source of the data to be copied to a target volume. The source and target volumes maintain their association through a copy pair. When the copy pair is removed, the source volume reverts back to a standard volume.
- **Target volume** – A standard volume becomes a target volume when it participates in a volume copy operation as the recipient of the data from a source volume. The source and target volumes maintain their association through a copy pair. When the copy pair is removed, the target volume reverts back to a standard volume.
- **Replicated volume** – A replicated volume is a volume that participates in a replication set. A replication set consists of two volumes; each is located on a separate array. After you create a replication set, the software ensures that the replicated volumes contain the same data on an ongoing basis.
- **Snapshot volume** – A snapshot volume is a point-in-time image of a standard volume. The management software creates a snapshot volume when you use the snapshot feature. The standard volume on which a snapshot is based is also known as the base or primary volume.

- **Reserve volume** – A snapshot reserve volume is automatically created when you create a snapshot. The reserve volume stores information about the data that has changed since the volume snapshot was created. When you delete a snapshot, the management software also deletes its associated reserve volume.

You can create up to 256 volumes on each virtual disk. During or after standard volume creation, you can map a host or host group to the volume in order to give the host or host group read/write privileges to the volume. Each host, including any host that is a member of a host group, must be assigned one or more initiators before the host or host group can be mapped to the volume. To see the current volumes, go to the Volume Summary page, as described in “Displaying Volume Information” on page 56.

There are a number of other things you can do with volumes, depending on their type, as described in TABLE 4-1.

TABLE 4-1 Additional Actions on Volumes

	Volume Type			
	Standard	Source	Target	Reserve
Map a volume to a host or host group	-	-	-	-
Unmap a volume from a host or host group	-	-	-	-
Create a volume snapshot	x	x	-	-
Recopy a volume copy	-	-	x	-
Copy a volume	x	x	x	-
Remove a copy pair	-	-	x	-
View performance statistics	x	x	x	-
Delete the volume	x	x	x	-

Planning Volumes

Creating a volume involves a number of tasks and decisions about a variety of elements in your storage configuration. Therefore, before running the New Volume wizard to create a new volume, you should plan your storage.

When you create a volume, be prepared to provide the following information:

- Volume name

Provide a unique name that identifies the volume.

- Volume capacity

Identify the capacity of the volume in megabytes, gigabytes, or terabytes.

- The storage pool for this volume

By default, the management software supplies a default storage pool. This pool uses the default storage profile, which implements RAID-5 storage characteristics that can be used in the most common storage environment. Other pools may have also been configured. Before you run the New Volume wizard, check the list of configured pools to see whether one of the pools has the storage characteristics you want. If a suitable pool does not exist, create a new pool using an existing or a new storage profile before you run the New Volume wizard.

- The way a virtual disk is selected

A volume can be created on a virtual disk as long as the RAID level, the number of disks, and the disk type (either FC or SATA) of the virtual disk matches the storage profile associated with the volume's pool. The virtual disk must also have enough capacity for the volume. You must choose the method of determining which virtual disk will be used to create the volume. The following options are available:

- Automatic – The management software automatically searches for and selects a virtual disk that matches the necessary criteria. If none are available, it creates a new virtual disk if enough space is available.
- Create Volume on an Existing Virtual Disk – You manually select the virtual disks on which to create the volume from the list of all available virtual disks. Be sure that the number of disks you select have enough capacity for the volume.
- Create a New Virtual Disk – You create a new virtual disk on which to create the volume. Be sure that the number of disks you select have enough capacity for the volume.

- Whether you want to map the volume now or later.

You can add the volume to an existing storage domain, including the default storage domain, or create a new one by mapping the volume to a host or host group. A storage domain is a logical entity used to partition storage that allows a host or host group to have read/write access to the volume. The default storage domain contains all hosts and host groups without explicit mappings and enables them to share access to all volumes that are not explicitly mapped. If you choose to map the volume later, the management software automatically includes it in the default storage domain.

Note: A host or host group will be available as a mapping option only if an initiator is associated with each individual host and each host included in a host group.

Managing Volumes

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

- “Displaying Volume Information” on page 56
- “Displaying Volume Performance Statistics” on page 57
- “Creating a Volume” on page 57
- “Modifying a Volume” on page 57
- “Copying Volume Information” on page 58
- “Mapping a Volume to a Host or Host Group” on page 58
- “Expanding Volume Capacity” on page 59
- “Unmapping a Volume From a Host or Host Group” on page 60
- “Changing Controller Ownership of a Volume” on page 60
- “Deleting a Volume” on page 61

Displaying Volume Information

You can display summary and detail information about existing storage volumes. You can also display information about the mapped hosts, mapped host groups, and snapshots that are associated with each volume.

To display information on volumes:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display volume information.
The navigation pane and the Volume Summary page for that array are displayed.
3. Click a volume name for detailed information on that volume.
The Volume Details page for the selected volume is displayed.
4. Go to Related Information and click any item for more information associated with the selected volume.
The Summary page for the selected item is displayed.

Displaying Volume Performance Statistics

To display information about volume performance:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see volume performance statistics.
The navigation pane and the Volume Summary page for that array are displayed.
3. Click View Performance Statistics.
The Performance Statistics Summary - Volumes page is displayed.

Creating a Volume

You must consider a number of factors and make a number of decisions before creating a volume. For information on planning the volume's storage characteristics, see "Planning Volumes" on page 54.

To create a volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to create a volume.
The navigation pane and the Volume Summary page for that array are displayed.
3. Click New.
The New Volume wizard is displayed.
4. Follow the steps in the wizard. Click the Help tab in the wizard for more information.

Modifying a Volume

To modify a volume's name or description:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to modify a volume.
The navigation pane and the Volume Summary page for that array are displayed.

3. Select the volume that you want to modify.
The Volume Details page for that volume is displayed.
4. Make the appropriate modifications, and click OK.
A message confirms that the volume was modified successfully.

Copying Volume Information

You can copy an existing volume to a target volume.

To copy an existing volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display volume information.
The navigation pane and the Volume Summary page for that array are displayed.
3. Click a volume name for detailed information on that volume.
The Volume Details page for the selected volume is displayed.
4. Click the Copy button.
5. The Copy Volume page is displayed.
6. Select a copy priority.
7. Select a target volume for the copy, and click OK.

Mapping a Volume to a Host or Host Group

To map a volume to a host or host group:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to map a volume.
The navigation pane and the Volume Summary page for that array are displayed.

3. Click the check box to the left of the volume to which you want to map a host or host group.

This enables the Map button.

4. Click Map.

The Map Volume page displays a list of available hosts and host groups. Use the filter to display only hosts or host groups.

5. Select the host or host group to which you want to map this volume, and click OK.

A message specifies that the selected volume was successfully mapped.

Expanding Volume Capacity

You cannot expand the capacity of a volume that has snapshots associated with it.

To expand the capacity of a volume:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array for which you want to expand volume capacity.

The navigation pane and the Volume Summary page for that array are displayed displayed.

3. Click the volume whose capacity you want to expand.

The Volume Details page is displayed.

4. Click Expand.

A message box displays, providing information about volume expansion.

5. Click OK.

The Expand Volume page displays the current capacity.

6. Specify the additional capacity you want, and click OK.

A message specifies that dynamic volume expansion is in progress.

Unmapping a Volume From a Host or Host Group

To unmap a volume from a host or host group:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to unmap a volume.

The navigation pane and the Volume Summary page for that array are displayed.

3. Click the volume that you want to unmap.

The Volume Details page is displayed.

4. Click Unmap.

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

Changing Controller Ownership of a Volume

The preferred controller owner for a volume is initially selected by the controller when a volume is created.

Under certain circumstances, ownership of a volume is automatically shifted so that the alternate controller becomes the current owner. For example, if the controller tray that is the preferred controller owner is being replaced or undergoing a firmware download, ownership of the volumes is automatically shifted to the other controller tray, and that controller becomes the current owner of the volumes. There can also be a forced failover from the preferred controller to the other controller because of I/O path errors.

Additionally, you can manually change the current owner of a volume to improve performance. For example, you might want to change the controller ownership of one or more volumes if you notice a disparity in the total input/output per second (IOPS) of the controllers such that the workload of one controller is heavy or is increasing over time while that of the other controller is lighter or more stable.

Note: You can change the controller ownership of a standard volume or a snapshot reserve volume. You cannot manually change the controller ownership of a snapshot volume because it inherits the controller owner of its associated base volume.

Use the Redistribute Volumes button to cause all volumes to return to their preferred controller owners.

To change a volume's controller owner:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to change controller ownership of a volume.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. Click the volume for which you want to change controller ownership.
The Volume Details page is displayed.
4. In the Owning Controller field, select either A or B.
5. Click OK.

Deleting a Volume

Before deleting a volume, consider the following:

- Deleting a volume that is mapped to a host or host group also causes the mappings to be deleted.
- Deleting a volume that has snapshots also causes the snapshots to be deleted.
- Deleting a volume that is part of a replication set also causes the replication set to be deleted. The remote volume remains intact, however.
- Deleting the only volume in a virtual disk causes the virtual disk to be deleted. If there are other volumes in the virtual disk, the storage space used by the deleted volume is converted to a free extent for future volume creation.

To delete a volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to delete a volume.
The navigation pane and the Volume Summary page for that array are displayed.
3. Select the check box for the volume that you want to delete.
This enables the Delete button.
4. Click Delete. The volume is removed from the Volume Summary table.

Configuring Volume Copies

This section describes volume copies. It contains the following subsections:

- “About Volume Copies” on page 62
- “Planning Volume Copies” on page 63
- “Managing Volume Copies” on page 65

About Volume Copies

A volume copy is a copy of data on one volume (called a source volume) written onto another volume (called a target volume) on the same storage array. A volume copy can be used to back up data, copy data from volumes that use small-capacity drives to volumes that use large-capacity drives, and restore snapshot data to the primary volume.

A source volume accepts host I/O and stores application data. The target volume maintains a copy of the data from the source volume.

When you create a volume copy, the management software creates a copy pair, which defines the association between the source volume and the target volume. When you no longer need a particular volume copy, you can remove the copy pair. Removing the copy pair dissociates the source volume from the target volume and causes the target and source volumes to revert their original volume types (such as standard volume or volume snapshot), enabling them to participate in other copy pairs in different roles. Removing a copy pair does not remove the data on the target volume.

The process of creating a volume copy is managed by the redundant array of independent disks (RAID) controllers and is transparent to host machines and applications. When the volume copy process starts, the controller of the source volume reads the data from the source volume and writes it to the target volume. The volume copy has a status of In-progress while the volume copy is being completed. Up to eight volume copies can have the status of In-progress at one time.

While a volume copy has a status of In-progress, the same controller must own both the source volume and the target volume. If different controllers own the source volume and target volume before the volume copy process is started, the management software automatically transfers ownership of the target volume to the controller that owns the source volume. When the volume copy process is completed or stopped, the management software restores ownership of the target volume to its original controller owner. Similarly, if ownership of the source volume is changed during the volume copy, ownership of the target volume is also changed.

The status of a volume copy can be one those described in TABLE 4-2.

TABLE 4-2 Volume Copy Status

Volume Copy Status	Description
Completed	The volume copy process finished successfully. <ul style="list-style-type: none">• The source volume is available for both read and write I/O activity.• The target volume is available for read I/O activity, unless the Read-Only attribute has been set to No.
In-progress	Data is being copied from the source volume to the target volume. A maximum of eight volume copies can be in progress at one time. <ul style="list-style-type: none">• The source volume is available for read I/O activity only.• The target volume is not available for read or write I/O activity.
Pending	The volume copy is waiting to be processed. If more than eight volume copies have been requested, subsequent volume copies enter a Pending state until one of the volume copies with a status of In-progress is completed. <ul style="list-style-type: none">• The source volume is available for read I/O activity only.• The target volume is not available for read or write I/O activity
Failed	The volume copy process failed. <ul style="list-style-type: none">• The source volume is available for read I/O activity only.• The target volume is not available for read or write I/O activity
Copy halted	The volume copy was stopped before all of the data on the source volume was copied to the target volume.

Planning Volume Copies

When you create a volume copy, be prepared to do the following:

- Select a source volume from the Volume Summary page or from the Snapshot Summary page.

A source volume can be any of the following volume types:

- Standard volume
- Snapshot
- Base volume of a snapshot (a volume of which you took a snapshot)
- Target volume

You can copy one source volume to several different target volumes.

- Select a target volume from the list of target volume candidates.

The target volume must have a capacity that is equal to or greater than the usable capacity of the source volume. A target volume must be one of the following:

- Standard volume
- Base volume of a Failed or Disabled volume snapshot.

Note: In order for a volume to be used as a target volume, its snapshots need to be either failed or disabled.

Caution: A volume copy will overwrite all data on the target volume and automatically make the target volume read-only to hosts. Ensure that you no longer need the data or have backed up the data on the target volume before starting a volume copy. After the volume copy process has finished, you can enable hosts to write to the target volume by changing the target volume's Read-Only attribute on the Volume Details page.

Because a target volume can have only one source volume, it can participate in one copy pair as a target. However, a target volume can also be a source volume for another volume copy, enabling you to make a volume copy of a volume copy.

- Set the copy priority for the volume copy.

During a volume copy, the storage array's resources may be diverted from processing I/O activity to completing a volume copy, which may affect the storage array's overall performance.

Several factors contribute to the storage array's performance, including I/O activity, volume redundant array of independent disks (RAID) level, volume configuration (number of drives and cache parameters), and volume type (volume snapshots may take more time to copy than standard volumes).

When you create a new volume copy, you will define the copy priority to determine how much controller processing time is allocated for the volume copy process and diverted from I/O activity.

There are five relative priority settings. The Highest priority rate supports the volume copy at the expense of I/O activity. The Lowest priority rate supports I/O activity at the expense of volume copy speed.

You can specify the copy priority before the volume copy process begins, while it is in progress, or after it has finished (in preparation for recopying the volume).

Managing Volume Copies

Before you can use the volume copy feature, you must enable it.

This section describes how to manage volume copies. It contains the following subsections:

- “Enabling the Volume Copy Feature” on page 65
- “Displaying Volume Copy Information” on page 66
- “Creating a Volume Copy” on page 66
- “Recopying a Volume Copy” on page 67
- “Changing the Copy Priority” on page 68
- “Removing a Copy Pair” on page 69

Enabling the Volume Copy Feature

To enable the volume copy feature:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to use the volume copy feature.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Add License.
The Add License page is displayed.
5. Select Volume Copying from the License Type menu.
6. Enter the version number and the key digest, and click OK.

Note: If you disable the volume copy feature, but volume copy pairs still exist, you can still remove the copy pair, start a copy using the existing copy pair, and change the setting of the read-only attribute for target volumes. However, you cannot create new volume copies.

Displaying Volume Copy Information

To determine which volumes are involved in a volume copy, see the Volume Details page. For the source volume, the Related Information section identifies associated target volumes. For the target volume, the Volume Details page identifies the associated source volume, the copy priority, and the target volume read-only status.

To display information on source volumes:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display volume copy information.
The navigation pane and the Volume Summary page for that array are displayed.
3. Click a source or target volume name for detailed information on that volume.
The Volume Details page for the selected volume is displayed.
4. Go to Related Information and click any item for more information associated with the selected volume.
5. Click Targets.
The Volume Copies Summary page lists the target volumes associated with selected volume.
6. Click a target volume name to display information about that volume.
The Volume Details page for the selected target volume is displayed.

Creating a Volume Copy

Before creating a volume copy, be sure that a suitable target volume exists on the storage array, or create a new target volume specifically for the volume copy. For information on planning volume copies, see “Planning Volume Copies” on page 63.

You can create a copy of a standard volume, a target volume, or a snapshot volume. For information about copying a snapshot volume, see “Copying a Volume Snapshot” on page 78.

To create a volume copy of a standard volume or a target volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array for which you want to create a volume copy.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. Click the name of the volume whose contents you want to copy to another volume. The volume you select must be either a standard volume, a snapshot volume, or a target volume.
The Volume Details page for that volume is displayed.
4. Click Copy.
5. When prompted to continue, click OK.
The Copy Volume page is displayed.
6. Select the copy priority.
The higher the priority you select, the more resources will be allocated to the volume copy operation at the expense of the storage array's performance.
7. Select the target volume you want from the Target Volumes list.
Select a target volume with a capacity similar to the usable capacity of the source volume to reduce the risk of having unusable space on the target volume after the volume copy is created.
8. Before starting the volume copy process:
 - a. Stop all I/O activity to the source and target volumes.
 - b. Unmount any file systems on the source and target volumes, if applicable.
9. Review the specified information on the Copy Volume page. If you are satisfied, click OK to start the volume copy.
A message confirms that the volume copy has successfully started.
10. After the volume copy process has finished:
 - a. Remount any file systems on the source volume and target volume, if applicable.
 - b. Enable I/O activity to the source volume and target volume.

Recopying a Volume Copy

You can recopy a volume copy for an existing copy pair. Recopying a volume copy is useful when you want to perform a scheduled, complete backup of the target volume that can then be copied to a tape drive for off-site storage.

Caution: Recopying a volume copy will overwrite all data on the target volume and automatically make the target volume read-only to hosts. Ensure that you no longer need the data or have backed up the data on the target volume before recopying a volume copy.

To recopy a volume copy:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to recopy a volume copy.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. Click the name of the target volume that you want to recopy.
The Volume Details page for that volume is displayed.
4. Stop all I/O activity to the source volume and target volume.
5. Unmount any file systems on the source volume and target volume, if applicable.
6. Click Recopy.
The management software recopies the source volume to the target volume and displays a confirmation message.
7. Remount any file systems on the source volume and target volume, if applicable.
8. Enable I/O activity to the source volume and target volume.

Changing the Copy Priority

To change the copy priority for a volume copy:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to change the copy priority of a volume copy.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. Click the name of the volume for which you want to change the copy priority.
The Volume Details page for the selected volume is displayed.

4. In the Copy Priority field, select the copy priority you want.

The higher the priority you select, the more resources will be allocated to the volume copy operation at the expense of the storage array's performance.

5. Click OK.

A confirmation message indicates that the change was successful.

Removing a Copy Pair

Removing a copy pair removes the association between a source volume and a target volume, enabling each to participate in another volume copy in a different role.

You cannot remove a copy pair while a volume copy using that copy pair is in progress.

To remove a copy pair:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array for which you want to remove a copy pair.

The navigation pane and the Volume Summary page for that array are displayed.

3. Click the name of the target volume in the copy pair that you want to remove.

The Volume Details page for that volume is displayed.

4. Click Related Information to ensure that a volume copy using the selected volume is not in progress.

5. Return to the top of the Details page, and click Remove Copy Pair.

Configuring Volume Snapshots

This section describes configuring and managing volume snapshots. It contains the following subsections:

- “About Volume Snapshots” on page 70
 - “Planning Volume Snapshots” on page 72
 - “Calculating Reserve Volume Capacity” on page 73
 - “Managing Volume Snapshots” on page 76
-

About Volume Snapshots

A snapshot is a copy of the data on a volume at the moment at which you create the snapshot. It is the logical equivalent of a complete physical copy, but you create it much more quickly than a physical copy, and it requires less disk space. Creating a snapshot causes the array controller tray to suspend I/O to the base volume while it creates a physical volume, called the reserve volume. The reserve volume stores information about the data that has changed since the snapshot was created. The capacity of the reserve volume is a configurable percentage of the base volume.

Snapshot is a premium feature that requires a license. You must enable the snapshot license before you can use the snapshot feature. For information on licenses, see “Managing Licenses” on page 42.

You work with a snapshot as you would with any other volume, with the exception that you cannot take a snapshot of a snapshot. Each snapshot can be accessed independently by other applications. A snapshot can be mounted on another server and used in the following ways:

- As an alternative backup method to reduce downtime for backup operations

Backing up snapshots, rather than online data, enables critical transactions to keep running during the backup process. Mount the snapshot to the backup server, and then back up the snapshot’s data to tape.

- For data analysis and testing of applications with actual, current data

Instead of working with operating data or interfering with critical transactions, use the most recent snapshot for analysis or testing of data.

- To restart applications

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

Note: A snapshot of data is not suitable for failure recovery. Continue to use offline backup methods to create full-volume backup copies.

Snapshots on the array are copy-on-write or dependent copies. In this type of snapshot, write operations to the primary volume causes the management software to copy the snapshot metadata and copy-on-write data to the reserve volume. Because the only blocks that are physically stored in the reserve volume are those that have changed since the time the snapshot was created, the snapshot uses less disk space than a full physical copy.

When a write operation occurs on the primary volume to a data block in which the data has not changed since the snapshot was created, the management software does the following:

- Copies the old data to the reserve volume
- Writes the new data to the primary volume
- Adds a record to the snapshot bitmap indicating the location of the new data

When a data host sends a read request to the snapshot, the management software checks whether the requested blocks have changed on the primary volume since the snapshot was created. If they have changed, the read request is satisfied from the data stored in the snapshot reserve volume. If blocks have not changed, the read request is satisfied from the primary volume. Snapshots can also accept write operations. Write operations to a snapshot are stored in the snapshot reserve volume.

The management software provides a warning message when the reserve volume nears the threshold, which is a configurable percentage of the full capacity of the snapshot reserve (the default is 50 percent). When the reserve volume threshold is met, the reserve volume's capacity can be expanded using the free capacity on the virtual disk.

As long as a snapshot is enabled, storage array performance is affected by the copy-on-write activity to the associated reserve volume. If a snapshot is no longer needed, you can stop the copy-on-write activity by either disabling or deleting the snapshot.

When a snapshot is disabled, it and its associated reserve volume still exist. When you need to create a different point-in-time image of the same primary volume, you can resnap the volume to reuse the disabled snapshot and its associated reserve volume. This takes less time than creating a new snapshot.

If you do not intend to re-create a snapshot, you can delete the snapshot instead of disabling it. When you delete a snapshot, the management software also deletes the associated reserve volume.

To see the current snapshots for a particular volume, go to the Related Information section of Snapshot Summary page for that volume, as described in "Displaying Volume Snapshot Information" on page 77.

Planning Volume Snapshots

Create a volume snapshot involves a number of tasks and decisions about a variety of factors. Therefore, before running the Create a Snapshot Volume wizard. The wizard, you should plan the following aspects of your snapshots:

- The name of the snapshot reserve volume

When you create a snapshot, you must provide a unique name for the snapshot that enables you to easily identify the primary volume.

Each snapshot has an associated reserve volume that stores information about the data that has changed since the snapshot was created. You must provide a unique name for the reserve volume that enables you to easily identify the snapshot to which it corresponds.

- The capacity of the reserve volume

To determine the appropriate capacity, you must calculate both the management overhead required and percentage of change you expect on the base volume. For more information, see “Calculating Reserve Volume Capacity” on page 73.

- The warning threshold

When you create a snapshot volume, you can specify the threshold at which the management software will generate messages to indicate the level of space left in the reserve volume. By default, the software generates a warning notification when data in the reserve volume reaches 50 percent of the available capacity. You can monitor the percentage of space used on the Snapshot Details page for the snapshot.

- The method used to handle snapshot failures

When you create a snapshot volume, you can determine how the management software will respond when the reserve volume for the snapshot becomes full. The management software can do either of the following:

- Fail the snapshot volume. In this case the snapshot becomes invalid, but the base volume continues to operate normally.
- Fail the base volume. In this case, attempts to write new data to the primary volume fail. This leaves the snapshot as a valid copy of the original base volume.

- The virtual disk selection method

A snapshot can be created on a virtual disk as long as the virtual disk has enough capacity for the snapshot.

The following options are available:

- Automatic – The management software automatically searches for and selects a virtual disk that matches the necessary criteria. If there are none, and enough space is available, it creates a new virtual disk.
- Create Volume on an Existing Virtual Disk – You manually select the virtual disks on which you want to create the volume from the list of all available virtual disks. Be sure that the number of disks you select have enough capacity for the volume.
- Create a New Virtual Disk – You create a new virtual disk on which to create the volume. Be sure that the virtual disk that you create has enough capacity for the volume.
- The snapshot mapping option

You can add the snapshot to an existing storage domain, including the default storage domain, or create a new storage domain by mapping the snapshot to a host or host group. A storage domain is a logical entity used to partition storage that allows a host or host group to have read/write access to the snapshot. The default storage domain contains all hosts and host groups without explicit mappings and enables them to share access to all snapshots that are not explicitly mapped.

During snapshot creation, you can choose between the following mapping options:

- Map Snapshot to One Host or Host Group - this option enables you to explicitly map the snapshot to a specific host or host group, or to include the snapshot in the default storage domain.
- Do Not Map this Snapshot - this option causes the management software to automatically include the snapshot in the default storage domain.

Note: A host or host group will be available as a mapping option only if an initiator is associated with each individual host and each host included in a host group.

Calculating Reserve Volume Capacity

When you create a snapshot, you specify the size of the snapshot reserve volume that will store snapshot data and any other data that is needed during the life of the snapshot. When prompted to specify the size of the snapshot reserve volume, you must enter a percentage of the size of the base volume, as long as that percentage does not translate to a size of less than 8 megabytes.

The capacity needed for the snapshot reserve volume varies, depending on the frequency and size of I/O writes to the base volume and how long you need to keep the snapshot volume. In general, choose a large capacity for the reserve volume if you intend to keep the snapshot volume for a long period of time or if you anticipate heavy I/O activity, which will cause a large percentage of data blocks to change on the base volume during the life of the snapshot volume. Use historical performance to monitor data or other operating system utilities to help you determine typical I/O activity on the base volume.

When the snapshot reserve volume reaches a specified capacity threshold, you are given a warning. You set this threshold when you create a snapshot volume. The default threshold level is 50 percent.

If you receive a warning and determine that the snapshot reserve volume is in danger of filling up before you have finished using the snapshot volume, you can increase its capacity by navigating to the Snapshot Details page and clicking Expand. If the snapshot reserve volume fills up before you have finished using the snapshot, the snapshot failure handling conditions specify the action that will be taken.

When you create a snapshot volume, you can allocate as large a snapshot reserve volume as you have space for.

Use the following information to determine the appropriate capacity of the snapshot reserve volume:

- A snapshot reserve volume cannot be smaller than 8 megabytes.
- The amount of write activity to the base volume after the snapshot volume has been created dictates how large the snapshot reserve volume needs to be. As the amount of write activity to the base volume increases, the number of original data blocks that need to be copied from the base volume to the snapshot reserve volume also increases.
- The estimated life expectancy of the snapshot volume contributes to determining the appropriate capacity of the snapshot reserve volume. If the snapshot volume is created and remains enabled for a long period of time, the snapshot reserve volume runs the risk of reaching its maximum capacity.
- The amount of management overhead required on the snapshot reserve volume for storage of snapshot volume data contributes to determining the appropriate capacity of the snapshot reserve volume. The amount of management overhead actually required is fairly small, and can be calculated with the simple formula that is presented later in this topic.
- There is not necessarily a one-to-one correlation between the number of data blocks that change on the base volume and the amount of data stored in the snapshot reserve volume. Depending on the location of data blocks that need to be copied, performance factors might dictate that the controller copy over a full

set of 32 blocks, even if only one set of blocks has changed. Consider this when determining the percentage of the base volume's capacity that can be copied to the snapshot reserve volume.

Use the following formula to calculate the amount of management overhead required to store snapshot data on the snapshot reserve volume.

$$192 \text{ Kbytes} + (x/2000)$$

where x is the capacity of the base volume in bytes.

Note: This formula is merely a guide. You should re-estimate the snapshot reserve volume capacity periodically.

The conversion process involves conversion from bytes to kilobytes and then to megabytes or gigabytes. For example, for a 5-gigabyte base volume, you would calculate the estimated snapshot reserve volume capacity as follows:

1. Convert the base volume's capacity to bytes.

When converted, 5 gigabytes equals 5,368,709,120 bytes.

2. Divide the base volume's capacity (in bytes) by 2000.

The result is 2,684,354.56 bytes.

3. Convert bytes to kilobytes.

The result is 2621.44 kilobytes.

4. Add 192 kilobytes to the results from Step 3.

$$192 \text{ Kbytes} + 2621.44 \text{ Kbytes} = 2813.44 \text{ Kbytes}$$

5. Convert the result from Step 4 to megabytes.

The resulting amount of management overhead required is 2.75 megabytes (or 0.002686 gigabytes).

Continuing in this example, suppose that you expect 30 percent of the data blocks on the base volume to change. To accurately calculate the snapshot reserve volume capacity, you must allow sufficient space for the snapshot reserve volume as well as for the management overhead.

To determine the snapshot reserve volume, calculate the percentage of change you expect on the base volume:

$$30 \text{ percent} \times 5 \text{ gigabytes} = 1.5 \text{ gigabytes}$$

To obtain the final estimated snapshot reserve volume capacity, add this number to the previously calculated amount of management overhead:

$$1.5 \text{ gigabytes} + 0.002686 \text{ gigabytes} = 1.502686 \text{ gigabytes}$$

In the Create Snapshot Volume Wizard: Specify Reserve Capacity dialog box, use the percentage (%) of base volume to specify the estimated capacity of the snapshot reserve volume.

When you create a snapshot, you will specify this snapshot reserve volume capacity as a percentage of the base volume. You can increase or decrease the percentage until the Snapshot Reserve Volume Capacity value matches the estimated capacity that you calculated. Some rounding up may be required.

Managing Volume Snapshots

Before you can use the snapshots feature, you must enable it.

This section describes managing volume snapshots. It contains the following subsections:

- “Enabling Volume Snapshots” on page 76
- “Displaying Volume Snapshot Information” on page 77
- “Creating a Volume Snapshot” on page 77
- “Resnapping a Volume Snapshot” on page 78
- “Copying a Volume Snapshot” on page 78
- “Mapping a Volume Snapshot to a Host or Host Group” on page 79
- “Expanding Snapshot Capacity” on page 80
- “Unmapping a Volume Snapshot” on page 80
- “Disabling a Volume Snapshot” on page 81
- “Deleting a Volume Snapshot” on page 82

Enabling Volume Snapshots

Before you can use the snapshots feature, you must enable it.

To enable the volume snapshots feature:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to enable snapshots.
The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Add License.
The Add License page is displayed.
5. Select Snapshot from the License Type menu.
6. Enter the version number and the digest, and click OK.

Displaying Volume Snapshot Information

You can display summary and detail information for existing snapshots.

To display information on snapshots:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display snapshot information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Click the snapshot name for detailed information on that snapshot.
The Snapshot Details page for the selected snapshot is displayed.

Creating a Volume Snapshot

You must consider a number of factors and make a number of decisions before creating a snapshot. For information on planning a snapshot, see “Planning Volume Snapshots” on page 72.

Note: You cannot create a snapshot of a target volume.

To create a volume snapshot:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to create a snapshot.
The navigation pane and the Volume Summary page are displayed.

3. Click the volume for which you want to create a snapshot.
The Volume Details page for that volume is displayed.
4. Click Snapshot.
The Create a Snapshot Volume wizard is launched.
5. Follow the steps in the wizard. Click the Help tab in the wizard for more information.

Resnapping a Volume Snapshot

To resnap a volume snapshot:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to resnap a snapshot.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Select the snapshot that you want to resnap.
This enables the Resnap button.
5. Click Resnap.
Upon completion of the new snapshot, a confirmation message is displayed.

Copying a Volume Snapshot

Before copying a snapshot, be sure that a suitable target volume exists on the storage array, or create a new target volume specifically for the snapshot.

Caution: Like any volume copy, a volume copy of a snapshot overwrites all data on the target volume and automatically makes the target volume read-only to data hosts. Ensure that you no longer need the data or have backed up the data on the target volume before starting a volume copy.

To copy a snapshot:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array for which you want to copy a snapshot.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Click the name of the snapshot that you want to copy.
The Snapshot Details page is displayed.
5. Click Copy.
The Copy Snapshot page is displayed.
6. Select a copy priority.
Valid values are Highest, High, Medium, Low, and Lowest. The higher the priority, the more resources will be allocated at the expense of the storage array's performance.
7. Select a target volume from the Target Volumes list.
Only valid target volumes with capacities equal to or greater than the capacity of the source snapshot are displayed.
Note: Selecting a target volume with a capacity similar to that of the source snapshot reduces the risk of there being unusable space on the target volume after the volume copy has been created.
8. Stop all I/O activity to the snapshot and target volume.
9. Unmount any file systems on the source volume and target volume, if applicable.
10. Review the specified information. If you are satisfied, click OK.
11. Remount any file systems on the source volume and target volume, if applicable.
12. Enable I/O activity to the snapshot volume and target volume.

Mapping a Volume Snapshot to a Host or Host Group

To map a volume snapshot to a host or host group:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to map a snapshot.
The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Click the snapshot to which you want to map the host or host group.
This enables the Map button.
5. Click Map.
The Map Volume page displays a list of available hosts and host groups. Use the filter to display only hosts and host groups.
6. Select the host or host group that you want to map to this volume, and click OK.
A message confirms that the snapshot was mapped successfully.

Expanding Snapshot Capacity

To expand the capacity of a snapshot volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to expand volume capacity.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Click the snapshot whose capacity you want to expand.
The Snapshot Details page is displayed.
5. Go to the Reserve Details section of the page, and click the Expand button.
The Expand Volume page displays the current capacity.
6. Specify the volume capacity you want, and click OK.
A message specifies that dynamic volume expansion is in progress.

Unmapping a Volume Snapshot

To unmap a volume snapshot:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array for which you want to unmap a snapshot.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Snapshots.
The Snapshot Summary page is displayed.
4. Select the snapshot that you want to unmap.
The Snapshot Details page is displayed.
5. Click Unmap.
You are prompted to confirm the unmapping.
6. Click OK.
When the unmapping is complete, a confirmation message is displayed.

Disabling a Volume Snapshot

Disabling a volume snapshot does not remove either the volume snapshot or its associated reserve volume.

When you disable a volume snapshot:

- You can re-enable the snapshot by selecting it from the Snapshot Summary page and clicking the Re-snap button.
- Only the specified snapshot is disabled. All other snapshots remain functional.

To disable a volume snapshot:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to disable a snapshot.
The navigation pane and the Volume Summary page are displayed.
3. Select the snapshot that you want to disable.
This enables the Disable button.
4. Click Disable.
You are prompted to confirm the disabling of the snapshot volume.
5. Click OK.
When the selected snapshot has been disabled, a confirmation message is displayed.

Deleting a Volume Snapshot

When you delete a volume snapshot, the corresponding reserve volume is also deleted, thus freeing up the capacity allocated to the reserve volume.

To delete a volume snapshot:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array for which you want to delete a snapshot.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Logical Storage > Snapshots.

The Snapshot Summary page is displayed.

4. Select the snapshot that you want to delete.

This enables the Delete button.

5. Click Delete.

You are prompted to confirm the deletion.

6. Click OK.

When the selected snapshot has been deleted, a confirmation message is displayed.

Configuring Data Replication

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

- “About Data Replication” on page 83
 - “Planning for Data Replication” on page 93
 - “Managing Data Replication” on page 97
-

About Data Replication

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

- “About the Replication Process” on page 83
- “About Replication Sets” on page 85
- “About the Consistency Group” on page 87
- “About Replication Links” on page 88
- “About Replication Set Properties” on page 88
- “About Primary and Secondary Role Reversal” on page 89
- “About Replication Modes” on page 91
- “Reference: Synchronous Versus Asynchronous Replication Modes” on page 92
- “About Data Replication Operations” on page 93

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 volumes between physically separate primary and secondary arrays in real time. The software is active while your applications access the volumes, and it continuously replicates the data between volumes.

As part of a disaster recovery and business continuance plan, the software enables you to keep up-to-date copies of critical data from the primary volume on the secondary volume. You can also rehearse your data recovery strategy to fail data over to the secondary volume. Later, you can write any data changes that occurred back to the primary volume.

The software replicates data from a primary volume to a secondary volume. The association between the primary and secondary volumes constitutes 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 replication mode. See “About Replication Modes” on page 91 for more information.

Note: Third-party applications can continue to write to the primary volume while it is replicating, but the secondary volume is read only.

The software transports data between the two arrays by means of synchronous or asynchronous replication mode, using a dedicated Fibre Channel (FC) connection (FC port 4 for the 6140 array, and host port 2 for the 6130 array). Both of the arrays with volumes participating in the replication set must be registered with Sun’s management software, and must be reachable through the out-of-band management network.

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 data replication strategies using multiple arrays 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 Client Solutions.

If there is a break in the network or if the secondary volume is unavailable, the software automatically switches to suspended mode, in which it ceases replication and tracks changes to the primary volume in a separate volume known as a replication repository. When communication is restored, the software uses the information in the replication repository volume to resynchronize the volumes and returns to replicating the data.

When replicating data in synchronous replication mode, 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 volume.

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 the consistency group. Each array supports only one consistency group and replication sets in the consistency group must use asynchronous replication mode. The consistency group enables you to manage several replication sets as one. By using the 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.

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

About Replication Sets

A replication set includes the following:

- A volume residing on an array and a reference to a volume residing on another, physically separate array. One array contains the primary volume, which copies the data, and the other array contains the secondary volume, which is the recipient of the data.
- The replication mode between both arrays: synchronous or asynchronous.
- The role that the volume plays within the replication set, either as a primary or as a secondary volume.

Up to 32 replication sets are supported per array. You can create and configure a replication set from either array. It is not necessary to set up the replication set properties on both arrays.

Note: Multi-hop and one-to-many data replication are not supported.

You can update the secondary volumes synchronously in real time or asynchronously using a store-and-forward technique. When the replication set is first created, a primary volume is first wholly 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.

When you activate the Sun StorageTek Data Replicator software premium feature on each array (see “Activating and Deactivating Data Replication” on page 99), two replication repository volumes are created on each array, one per controller. The controller stores replication information in the replication repository volume, including information about write operations to the destination volume in the replication set that are not yet completed. The controller can then use this information to recover from controller resets or array outages, by copying only the blocks that have changed since the reset or outage occurred.

FIGURE 4-1 shows the relationship between the two arrays and their corresponding replication sets (for simplicity, only one controller/replication repository volume on each array is shown).

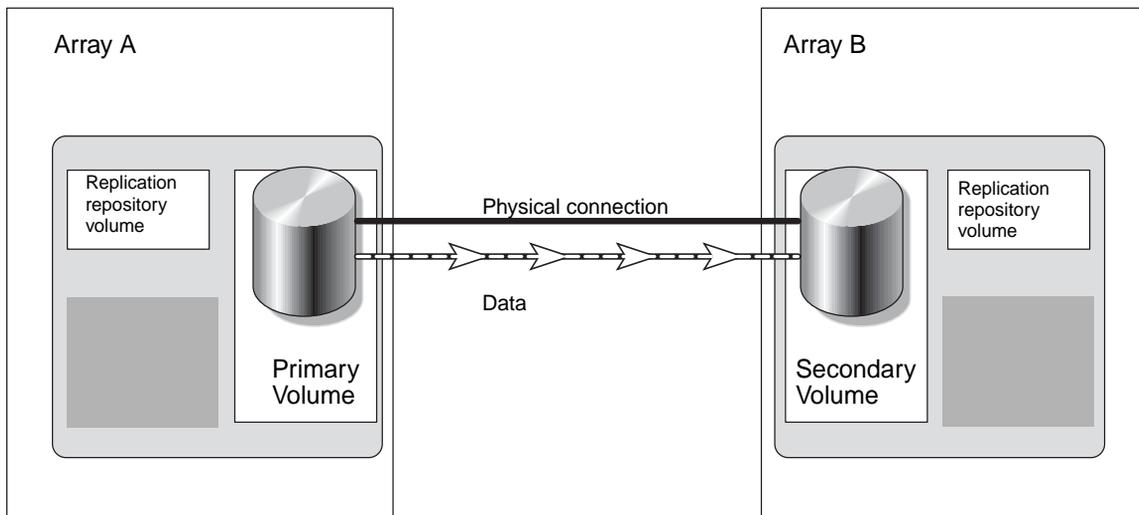


FIGURE 4-1 Relationship Between Arrays and Their Corresponding Replication Sets

After you create a replication set, you can modify its properties (see “About Replication Set Properties” on page 88 for more information). You can also perform volume operations, such as the following:

- Extend replicated volumes, by adding storage first to the secondary volume and then to the primary volume.
- Create snapshots.

You can create a snapshot of either the primary or the secondary volume within a replication set at any time.

- Make volume copies.

A primary volume in a replication set can be a source volume or a target volume in a volume copy. Volume copies are not allowed on secondary volumes. If a copy of a secondary volume is required, perform a role reversal to change the secondary volume to a primary volume. If a role reversal is initiated during a volume copy in progress, the copy will fail and you cannot restart it. For more on role reversal, see “About Primary and Secondary Role Reversal” on page 89.

- Change volume mappings.

You can also delete a replication set. This removes the association between the primary and secondary volumes, and the volumes revert to independent volumes.

Note: You cannot replicate a volume that is already in a replication set.

If the primary volume becomes unavailable, the secondary volume assumes 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.

About the Consistency Group

The consistency group is a collection of replication sets that have the same role and that only use asynchronous replication mode. The purpose of the consistency group is to ensure write order consistency across multiple volumes. Only one consistency group is supported per array.

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

When you include replication sets in the 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. 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 include a replication set in the consistency group, consider the guidelines in the TABLE 4-3.

TABLE 4-3 Consistency Group Guidelines

Regarding	Note
Write order	If you need to preserve write order consistency for a set of volumes, include the replication sets that comprise the volumes in the consistency group. You can add an existing replication set to the consistency group from the replication set's Replication Set Details page.
Volumes	All primary volumes in the consistency group must reside on the same primary array. You cannot preserve write ordering when you have primary volumes originating on different arrays.

TABLE 4-3 Consistency Group Guidelines (*Continued*)

Regarding	Note
Replication sets	Replication sets in the consistency group must all use asynchronous replication mode. The consistency group supports up to 32 replication sets.

A best practice is to plan the consistency group in advance and include the replication set in the consistency group at the same time that you create the replication set. To do this, select the Add to Write Consistency Group check box when you are specifying replication properties in the Create Replication Set wizard.

To add an existing replication set with asynchronous replication mode to the consistency group, select the Consistency Group check box on the Replication Set Details page for the replication set, then click OK to save.

About Replication Links

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

For the Sun StorageTek 6130 and 6140 arrays, you must use Fibre Channel (FC) port 2 and 4, respectively, on both arrays to establish the replication link between the arrays. Additionally, you must configure any FC switches that you use to make the connection to the array for long-distance operations and apply zoning practices. For more information on configurations and switch zoning for data replication, see the Getting Started guide for your array.

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 data replication strategies using multiple arrays 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 Client Solutions.

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

About Replication Set Properties

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

- Specify a secondary volume. When you create the replication set, the local volume assumes the primary role by default, and you are prompted to choose the secondary volume. Once the replication set is created, you can change the role of a volume on the Replication Set Details page as needed without suspending replication. See “Reversing Roles” on page 102 for more information.
- Set the replication mode to synchronous or asynchronous. See “About Replication Modes” on page 91 for more information. If the mode is asynchronous, you have the option of including the replication set in the consistency group. See “About the Consistency Group” on page 87 for more information.
- Set the priority of the synchronization rate relative to I/O activity to highest, high, medium, low, or lowest.
- If appropriate for your environment, set the resynchronization method in the Create Replication Set wizard to Automatic.

Automatic resynchronization is an alternative to manual synchronization (you manually synchronize by clicking the Resume button on the Replication Set Details page). The Automatic resynchronization option supports both replication sets and the consistency group. If you enable this option, the software synchronizes the volumes on both arrays and resumes replication as soon as possible.

For example, if a network link fails and causes the software to cease replication, resynchronization will occur when the link is restored. If the replication set is a member of the consistency group and it becomes unsynchronized (replication is suspended) due to a link failure or other problem, all replication sets in the consistency group with primary volumes on the local array will become unsynchronized. When the problem is resolved, all of the replication sets will automatically become resynchronized.

However, the software will not perform an automatic resynchronization if you manually suspended a replication set (or all of the replication sets in the consistency group, if the replication set is a member of the group). In this case, you will need to click the Resume button to resynchronize the replication sets.

You can enable automatic resynchronization for an existing replication set by selecting the Auto-synchronize check box on the Replication Set Details page. This check box is available only if the local volume is the primary volume.

About Primary and Secondary Role Reversal

A role reversal promotes the secondary volume to the primary volume within the replication set, and demotes the primary volume to the secondary volume.

Note: If you change the role of a volume in a replication set that is a member of the consistency group, the replication set will become a member of the consistency group on the array that hosts the newly promoted primary volume.

Caution: If I/O operations to the primary and secondary volumes occur during a role reversal, data is lost when you initiate a synchronization operation. 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.

Secondary to Primary

You usually promote a secondary volume to a primary volume in a replication set when a catastrophic failure has occurred on the array that contains the primary volume. You need to promote the secondary volume so that host applications that are mapped to the primary volume can still access data and so that business operations can continue. If the replication set is operating normally with a viable link, promoting the secondary volume to the primary volume automatically demotes the primary volume to the secondary volume.

When the secondary volume becomes a primary volume, any hosts that are mapped to the volume through a volume-to-LUN mapping will now be able to read or write to the volume. If a communication problem between the secondary and primary sites prevents the demotion of the primary volume, an error message is displayed. However, you are given the opportunity to proceed with the promotion of the secondary volume, even though this will lead to a dual-primary condition.

To promote a secondary volume to a primary volume, see “Reversing Roles” on page 102.

Primary to Secondary

You can demote a primary volume to a secondary role during normal operating conditions. If the replication set is operating normally with a viable link, demoting the primary volume to the secondary volume automatically promotes the secondary volume to the primary volume.

When the secondary volume becomes a primary volume, any host that accesses the primary volume through a volume-to-LUN mapping will no longer be able to write to the volume. When the primary volume becomes a secondary volume, only remote writes initiated by the primary controller will be written to the volume. If a communication problem between the primary and secondary sites prevents the promotion of the secondary volume, an error message is displayed. However, you are given the opportunity to proceed with the demotion of the primary volume, even though this will lead to a dual-secondary condition.

To demote a primary volume to a secondary volume, see “Reversing Roles” on page 102.

Note: To correct a dual-primary or dual-secondary condition, you must delete the replication set and then create a new one.

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 secondary volume has been updated. Synchronous replication forces the software to wait until the primary volume receives an acknowledgment of the receipt of the data from the secondary volume before returning to the application. This mode offers the best chance of full data recovery from the secondary volume, at the expense of host I/O performance.
- **Asynchronous mode** – In asynchronous mode replication, data is written to the primary volume and confirmed as complete before the secondary volume has been updated. Asynchronous replication enables the data replication software to return to the primary volume as soon as the write operation has been completed on the primary volume. The data is then copied to the secondary volume. This mode provides faster I/O performance, but does not guarantee that the copy to the secondary volume was completed before processing the next write request.

You can change the replication mode at any time during the life of a replication set, and you do not need to suspend replication before a mode change.

Note: If a replication set is a member of the consistency group, it must, by definition, use asynchronous mode replication. If you change the replication set to synchronous mode replication, it will no longer be part of the consistency group and will not affect the properties of the remaining replication sets in the group in any way.

Reference: Synchronous Versus Asynchronous Replication Modes

When you choose a replication mode, consider the characteristics of each, as described in TABLE 4-4.

TABLE 4-4 Synchronous and Asynchronous Mode Comparison

Synchronous Replication	Asynchronous Replication
A write operation to the primary volume is not considered complete until the user data is written to the secondary volume.	The write operation is considered complete as soon as the user data is written to the primary volume. The data is then copied to the secondary volume.
The response time depends on the network. The network latency must be low enough that your 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. The bandwidth of the network link must be able to handle the write traffic generated during the application's average write period.
Provides high data availability at the expense of longer latency across the wide area network.	Offers higher performance and availability in exchange for delayed synchronization of all copies of data.
Might result in an increase in 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).	Can potentially result in a longer response time than synchronous replication.
Because every data change is replicated to the secondary volume in real-time, the secondary volume is write-order consistent with the primary volume.	The secondary volume lags behind the primary volume, depending on the latency and the number of writes.
Write ordering across volumes is preserved at the secondary array.	If write ordering is required across volumes at the secondary array, write consistency must be enabled for the volumes at the primary array.
This mode is meant for volumes that have zero tolerance of data loss at the secondary array.	This mode is meant for volumes that have some level of tolerance for data loss at the secondary array.

About Data Replication Operations

You can choose one of the following data replication operations from the Replication Set Details page:

- **Suspend** – This operation temporarily stops replication of a replication set (or all of the replication sets in the consistency group with primary volumes on the primary array). While in suspended mode, the software logs any changes to the primary volume in the replication repository volume, and all writes to the secondary volume are blocked. No replication occurs.

You can use a suspend operation to save on telecommunications or connection costs. However, you risk data loss. If replication is suspended and then the primary volume fails, you do not have the data at the secondary volume that was written to the primary volume.

- **Resume** – This operation can occur only after suspension of a replication set (or of all of the replication sets in the consistency group with primary volumes on the primary array). At a later time, when the link is re-established or the problem resolved, the software uses the information in the replication repository volume to resynchronize the volumes. After the volumes are resynchronized, replication resumes.

When you choose to resume replication, consider the following:

- You can resume replication only if the local volume is the primary volume in the replication set (or the replication sets in the consistency group have their primary volumes on the local array).
- It is possible that large quantities of I/O will occur over the replication links as the volumes are brought back into synchronization. Be aware that a resynchronization operation may affect bandwidth.

Planning for Data Replication

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

- “General Planning Considerations” on page 94
- “Planning to Create Replication Sets” on page 95

General Planning Considerations

Replicating data and modifying replication properties may require 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 volume to become? This determines the replication mode and how often you should back up the data. 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.
- **Data loss** – 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 data on the primary volume, the data on the secondary volume is vulnerable during synchronization because write order is not preserved. Therefore, to ensure a high level of data integrity on both volumes during normal operations or data recovery, back up the data on both volumes before you perform a synchronization operation so that you always have a consistent copy of your data. If a failure occurs, the backup provides a “known good” copy from which your data can be restored.
- **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 volumes. 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 data 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 volumes.

Planning to Create Replication Sets

Before you can replicate data to a secondary volume, use the Create Replication Set wizard to create a replication set. Before you use the wizard, you must do the following:

- The storage domain definitions for the primary and secondary arrays are independent of each other. If these definitions are put into place while a volume is in a secondary role, it will reduce the administrative effort associated with site recovery if it becomes necessary to promote the volume to a primary role.
 - The size of the secondary volume must be equal to or greater than the size of the corresponding primary volume.
 - Primary and secondary volumes do not need to have the same redundant array of independent disks (RAID) level for replication.
 - The synchronization process overwrites all data on the secondary volume and makes it read-only. If you need to save any data on the secondary volume, back it up before creating the replication set.
 - You cannot use a replicated volume or a snapshot volume to create a replication set.
 - Any host mapped to a volume will no longer have write access to it once it becomes a secondary volume in a replication set. For this reason, avoid using mapped volumes as secondary volumes. However, any defined mappings will remain after the replication set is created, and any mapped host can resume writing to the volume if it is ever promoted to a primary volume or the replication set is deleted.
 - A volume can be the secondary volume for only one replication set.
- Make sure that you have configured the volumes for data replication on both arrays and that a viable Fibre Channel (FC) link exists between the arrays. Configure both the primary and secondary volumes as you would any other volume. The capacity of the secondary volume must be equal to or greater than that of the primary volume.
- Make sure that the secondary array is connected to the local primary array by means of an existing FC link. Data replication uses a dedicated FC port (FC port 2 on the 6130 array; FC port 4 on the 6140 array).
- Configure the FC switches that provide the connection between arrays for long-distance operations. See the FC switch vendor's documentation for information about operating over long distances. For more information on configurations and switch zoning for data replication, see the Getting Started Guide for your array.
- Make sure that both of the arrays with volumes participating in the replication set are registered with Sun's management software, and are reachable through the out-of-band management network.
- Record the names of the secondary array and secondary volume. You can find these names on the Volume pages for the secondary volume.

- Back up the data on both volumes so that you can restore it easily if there is a problem.
- Ensure that the Sun StorageTek Data Replicator software premium feature is enabled and activated on both arrays participating in data replication, as described in “Enabling Data Replication” on page 98 and “Activating and Deactivating Data Replication” on page 99.
- Define the characteristics of the replication set:
 - **Write order consistency** – Determine whether you need to preserve write order consistency across volumes and manage the volumes as a group. If so, you will include the replication set in the consistency group by selecting the Add to Write Consistency Group check box. See “About the Consistency Group” on page 87 for more information about the consistency group.
 - **Roles** – Determine which array and volume will assume the primary role. By default, the array and volume where you are running the Create Replication Wizard are assigned the primary role. For information about changing roles for existing replication sets, see “About Primary and Secondary Role Reversal” on page 89.
 - **Replication mode** – Determine whether synchronous or asynchronous data replication mode is the best method of data communication for your application. Consider network latency, bandwidth, and security when you choose the path through the storage area network (SAN) that you will use for data replication.

Note: The full synchronization that occurs when you first create the replication set is the most time consuming data replication operation. For this reason, avoid creating a new replication set across a limited-bandwidth link.

For more information about replication modes, see “About Replication Modes” on page 91.
 - **Synchronization priority** – 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 slower synchronization rate to ensure quicker I/O activity.
 - **Synchronization method** – Determine whether you want the system to automatically synchronize the volumes on both arrays whenever there is a working link, or whether you want to initiate synchronization manually. For more information about the Resynchronization Method option, see “About Data Replication Operations” on page 93 and “About Replication Set Properties” on page 88.

Managing Data Replication

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

- “About Data Replication Status” on page 97
- “Enabling Data Replication” on page 98
- “Activating and Deactivating Data Replication” on page 99
- “Disabling Data Replication” on page 100
- “Displaying Replication Set Information” on page 100
- “Creating Replication Sets” on page 101
- “Deleting Replication Sets” on page 102
- “Reversing Roles” on page 102
- “Changing Replication Modes” on page 103
- “Suspending and Resuming Data Replication” on page 104
- “Testing Replication Set Links” on page 105
- “Troubleshooting Data Replication” on page 105

About Data Replication Status

The data replication status is managed independently for the primary array and the secondary array. There are four possible data replication statuses, described in TABLE 4-5.

TABLE 4-5 Data Replication Status

Status	Description
Disabled/Deactivated	No data replication functions can be performed. The Sun StorageTek Data Replicator software premium feature is not available until it is enabled and activated. To enable data replication, see “Enabling Data Replication” on page 98. To activate data replication, see “Activating and Deactivating Data Replication” on page 99.
Disabled/Activated	Data replication is disabled, preventing new replication sets from being created. However, existing replication sets can be maintained with all functions of the Sun StorageTek Data Replicator software premium feature. To enable data replication, see “Enabling Data Replication” on page 98.

TABLE 4-5 Data Replication Status (*Continued*)

Status	Description
Enabled/Deactivated	Data replication is enabled, but not activated. Until data replication is activated, you cannot use any of the data replication functions. To activate data replication, see “Activating and Deactivating Data Replication” on page 99.
Enabled/Activated	Data replication is enabled and activated. Replication sets can be created and maintained with volumes on this array and any other array on which data replication is enabled and activated. To create and manage replication sets, see “Managing Data Replication” on page 97.

To determine whether data replication is enabled or disabled for an array, select the array, and then choose Administration > Licensing in the navigation pane to view the Licensable Feature Summary page.

To determine whether the data replicator software has been activated click Replication Sets, in the Available Features section on the Licenseable Feature Summary page, to view the Licenseable Feature Details - Replication Sets page.

Enabling Data Replication

Installing the license for the Sun StorageTek Data Replicator software premium feature on an array enables data replication for that array only. Since two arrays participate in a replication set, you must install a license on both arrays that you plan to have participate in a replication set.

Note: The 6130 array dedicates Fibre Channel (FC) port 2, and the 6140 array dedicated host port 4 on each controller for use with the Sun StorageTek Data Replicator software premium feature. Before enabling data replication on an array, you must ensure that FC port 2 on each controller on the 6130 array and FC port 4 on each controller on the 6140 array is not currently in use. If it is in use, you must move all connections from FC port 2 to FC port 1 for the 6130 array, and from FC port 4 to FC port 3, 2, or 1 for the 6140 array.

To enable data replication on an array:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to enable data replication.

The navigation pane and the Volume Summary page for that array are displayed displayed.

3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Add License.
The Add License page is displayed.
5. Select Sun StorageTek Data Replicator Software from the License Type menu.
6. Enter the version number and the key digest, and click OK.

Activating and Deactivating Data Replication

Activating the Sun StorageTek Data Replicator software premium feature prepares the array to create and configure replication sets. After data replication is activated, the secondary ports for each of the array's controllers are reserved and dedicated to data replication. In addition, a replication repository volume is automatically created for each controller in the array.

Note: The replication repository volumes require 256 megabytes of available capacity on an array. The two replication repository volumes are created with this capacity, one for each controller.

If no replication sets exist and the Sun StorageTek Data Replicator software premium feature is no longer required, you can deactivate data replication in order to re-establish normal use of dedicated ports on both storage arrays and delete both replication repository volumes.

Note: You must delete all replication sets before you can deactivate the premium feature.

To activate or deactivate the Sun StorageTek Data Replicator software premium feature:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array containing the primary volume in the data replication set.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Replication Sets.
The Licenseable Feature Details - Replication Sets page is displayed.

5. Click Activate or Deactivate, as appropriate.
A confirmation dialog box indicates success or failure.

Disabling Data Replication

When data replication is in the disabled/activated state, previously existing replication sets can still be maintained and managed; however, new data replication sets cannot be created. When in the disabled/deactivated state, no data replication activity can occur.

To disable data replication:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to locate the primary volume in the data replication set.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click the check box to the left of Replication Sets.
This enables the Disable button.
5. Click Disable.

Displaying Replication Set Information

To display information about the replication sets in an array:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the name of the array on which the replication set resides.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Replication Sets tab.
The Replication Set Summary page is displayed.

4. Click the name of a replication set.

The Replication Set Details page for the selected replication set is displayed.

Creating Replication Sets

You must do three things before creating a replication set:

- Consider a number of factors and make a number of decisions. For information on planning a replication set, see “Planning for Data Replication” on page 93.
- Stop all I/O activity and unmount any file systems on the secondary volume. Do this just before creating the replication set.
- Log in to the system using the `storage` user role.

The Create Replication Set wizard enables you to create a replication set, either standalone or as part of the consistency group. To create

To create a replication set:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the name of the array containing the primary volume that you want to replicate to the secondary volume.

The navigation pane and the Volume Summary page are displayed.

3. Click the name of the primary volume that you want to replicate to the secondary volume.

The Volume Details page for the selected volume is displayed.

Note: You cannot replicate a volume that is already in a replication set.

4. Click Replicate.

The Create Replication Set wizard is displayed.

5. Follow the steps in the wizard. The Create Replication Set wizard also allows you to include the new replication set in the consistency group, if desired.

When creating the replication set, the system copies all data from the primary volume to the secondary volume, overwriting any existing data on the secondary volume. If replication is suspended, either manually or due to a system or communication problem, and then resumed, only the differences in data between volumes are copied.

Note: An alternative method of creating a replication set is to go to the Replication Set Summary page and click the New button. In this case, an additional step in the wizard prompts you to filter and select the primary volume from the current array.

Deleting Replication Sets

When you delete a 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 data stored on the volumes is not affected, and the volumes can remain mapped. You can delete multiple replication sets simultaneously.

To remove the replication relationship between volumes:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the name of the array on which either volume in the replication set resides.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Replication Sets.
The Replication Set Summary page is displayed.
4. Select the replication set that you want to delete, and click Delete.
A confirmation message is displayed.
5. Click OK.
The replication set is deleted.

Reversing Roles

You can reverse the roles of the primary and secondary volumes in a replication set in the event of a failure of the primary volume. The secondary volume assumes the role of the primary volume so that the application software now accesses the secondary volume directly while the primary volume failure is corrected.

You can perform the role reversal from either volume in the replication set. For example, when you promote the secondary volume to a primary role, the existing primary volume is automatically demoted to a secondary role (unless the system cannot communicate with the existing primary volume).

Note: If you change the role of a volume in a replication set that is a member of the consistency group, the replication set will become a member of the consistency group on the array that hosts the newly promoted primary volume.

To reverse the role of volumes within a replication set:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the name of the array containing the volume in the replication set whose role you want to reverse.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Logical Storage > Replication Sets.
The Replication Set Summary page is displayed.

4. Click the name of the replication set that includes the volume.
The Replication Set Details page is displayed.

5. Click Role to Secondary or Role to Primary, as appropriate.
A confirmation message is displayed.

6. Click OK.

The roles of the volumes are now reversed.

Changing Replication Modes

You must consider a number of factors and make a number of decisions before changing the replication mode of a replication set. For information on planning replication modes, see “About Replication Modes” on page 91.

To change the replication mode of a replication set:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the name of the array containing the replication set whose replication mode you want to change.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Logical Storage > Replication Sets.
The Replication Set Summary page is displayed.

4. Click the name of the replication set whose replication mode you want to change.
The Replication Set Details page is displayed.

5. Select Asynchronous or Synchronous, as appropriate, from the mode drop-down list.

If you select Asynchronous, write order consistency is disabled by default. To enable write order consistency for all replication sets using asynchronous mode, select the Consistency Group check box.

6. Click OK to save the changes.

Suspending and Resuming Data Replication

To suspend or resume data replication in an existing replication set:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the name of the array containing the replication set for which you want to suspend or resume replication.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Logical Storage > Replication Sets.

The Replication Set Summary page is displayed.

4. Click the name of the replication set for which you want to suspend or resume replication.

The Replication Set Details page is displayed.

5. Do one of the following:

- If you want to suspend replication and track changes between the volumes, click Suspend.

Note: If the replication set is already in a Suspended, Unsynchronized, or Failed/Suspended state, only the Resume button is available.

Note: Suspending a replication set will stop the coordination of data between the primary and the secondary volume. Any data that is written to the primary volume will be tracked while the replication set is suspended and will automatically be written to the secondary volume when replication is resumed. A full synchronization will not be required.

- If you want to resume replication and copy only the data changes, not the entire contents of the volume, click Resume.

Note: Any data that is written to the primary volume will be tracked while the replication set is suspended and will automatically be written to the secondary volume when replication is resumed. A full synchronization will not be required.

6. When prompted to confirm the selected action, click OK.

Note: If you are suspending or resuming replication for a replication set that is part of the consistency group, all other replication sets in the group with primary volumes on the primary array will also be suspended or resumed.

Testing Replication Set Links

You can test communication between volumes in a replication set by clicking the Test Communication button on the Replication Set Details page. If a viable link exists between primary and secondary volumes, a message displays indicating that communication between the primary and secondary volume is normal. If there is a problem with the link, a message displays details about the communication problem.

Troubleshooting Data Replication

If data replication does not occur, make sure that:

- There is a working Fibre Channel connection on port 2 on both the primary and secondary 6130 array controllers or on port 4 on both the primary and secondary 6140 array controllers. To test link communication, see “Testing Replication Set Links” on page 105.
- You selected the correct secondary array and volume names in the Create Replication Set wizard. If you did not, you must delete the replication set and create a new one with the correct secondary array and volume names.

TABLE 4-6 describes some of the other problems that you may encounter when using data replication, possible causes, and resolution.

TABLE 4-6 Troubleshooting Data Replication Table

Symptom	Cause	Resolution
After a replication set is deleted on the primary array, the secondary array shows that the replication set still exists.	During the deletion operation, there was a link failure between primary and secondary arrays.	Delete the replication set on the secondary array.
After a replication set is deleted on the secondary array, the primary array shows that the replication set still exists.	During the deletion operation, there was a link failure between primary and secondary arrays.	Delete the replication set on the primary array.
The Synchronization Progress field shows an Unsynchronized state for the primary volume.	The data replicated between the primary and secondary volumes is no longer identical. Link errors, a failed primary or secondary volume, or a dual-primary or dual-secondary condition (see “About Primary and Secondary Role Reversal” on page 89) can all cause an Unsynchronized state.	If the controller owner of the primary volume can communicate with the controller owner of the secondary volume and the volume is online and replicating, and the Auto-synchronize option is selected, a full synchronization automatically takes place. Otherwise, click the Resume button to start full synchronization.

Configuring Virtual Disks

This section describes virtual disks. It contains the following subtopics:

- “About Virtual Disks” on page 106
- “Displaying Virtual Disk Information” on page 107
- “Expanding a Virtual Disk” on page 107
- “Defragmenting a Virtual Disk” on page 108

About Virtual Disks

Virtual disks are created and removed indirectly through the process of creating or deleting volumes or snapshots. The disk drives that participate in the virtual disk must all be of the same type, either Serial Attached Technology Advancement (SATA) or Fibre Channel (FC). 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 – that is, the array’s RAID level.

In the array, disk drives within a tray are grouped together into RAID sets, also called virtual disks, according to RAID level. The array supports the following levels:

- **RAID-0** – Stripes data across multiple disks, but without redundancy. This improves performance but does not deliver fault tolerance.
- **RAID-1** – Mirrors a disk so that all data is copied to a separate disk.
- **RAID-1+0** – Combines disk mirroring (RAID-1) with disk striping (RAID-0) to combine high performance and high levels of redundancy. RAID-0’s data striping is a cost-effective way to create high levels of performance in a disk array, and having multiple copies of data is the best way to create redundancy.

For the 6130 and 6140 arrays, to combine disk mirroring with disk striping, configure RAID-1 with more than 2 drives. The firmware automatically creates a RAID 1+0 virtual disk.

- **RAID-3** – Stripes data at the byte level across multiple disks while writing the stripe parity to a parity disk. Provides high throughput for a single streamed file. Checks parity on reads.
- **RAID-5** – Stripes data at the byte level, and provides stripe error correction (parity checking) information. 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, which improves the efficiency of available storage space.

Displaying Virtual Disk Information

You can display summary and detail information on existing virtual disks. You can also display summary information about the disk drives and volumes associated with each virtual disk.

To display information on virtual disks:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed
2. Click the array for which you want to display virtual disk information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Virtual Disks.
The Virtual Disk Summary page is displayed.
4. Click a virtual disk name for detailed information on that virtual disk.
The Virtual Disk Details page for the selected virtual disk is displayed.
5. Go to Related Information and click any item for more information associated with the selected virtual disk.
The Summary page for the selected item is displayed.

Expanding a Virtual Disk

You can expand the size of a virtual disk if there are enough unassigned disks and the virtual disk is online and ready.

Note: When you expand a virtual disk, the disk drives that participate in the virtual disk must all be of the same type, either Serial Attached Technology Advancement (SATA) or Fibre Channel (FC).

To expand a virtual disk:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed
2. Click the array for which you want to expand the virtual disk size.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Virtual Disks.
The Virtual Disk Summary page is displayed.
4. Click the name of the virtual disk you want to expand.
The Virtual Disk Details page for that virtual disk is displayed.
5. Click Expand to see the disks available for the expansion of this virtual disk.
The Expand Virtual Disk page is displayed.
6. Select the disks you want to include in the virtual disk to expand its capacity, and click OK.
The Virtual Disk Details page is refreshed, and a message confirms the virtual disk expansion.

Defragmenting a Virtual Disk

To defragment a virtual disk:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to defragment a virtual disk.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Virtual Disks.
The Virtual Disk Summary page is displayed.
4. Click the name of the virtual disk that you want to defragment.
The Virtual Disk Details page for that virtual disk is displayed.

5. Click Defragment.

The Virtual Disk Details page is refreshed, and a message confirms that the defragmentation process has successfully started.

Configuring Storage Pools

This section describes storage pools. It contains the following subsections:

- “About Storage Pools” on page 110
- “Displaying Pool Information” on page 110
- “Creating a Storage Pool” on page 111
- “Modifying a Storage Pool” on page 112
- “Deleting a Storage Pool” on page 112

About Storage Pools

A storage environment can be divided into storage pools. Each pool is associated to a profile that allows the storage pool to meet specific I/O requirements. Each array has a default pool that uses the default profile which implements RAID-5 storage characteristics.

The default pool satisfies most common storage requirements. The array also provides a set of storage profiles that satisfy certain I/O requirements which are optimal for the type of application to which they refer. If none of the factory profiles are suitable for your site’s needs, you can create a custom storage profile. When you create a new storage pool, you assign a specific profile to it.

To see the current storage pools, go to the Storage Pool Summary page, as described in “Displaying Pool Information” on page 110.

Note: Removing a storage pool destroys all stored data in the pool and removes all volumes that are members of the pool. The data can be restored from backup after new storage pools are added, but it is far easier to avoid the difficulty in the first place.

Displaying Pool Information

You can display summary and detail information on existing storage pools. You can also display summary information about the virtual disks and volumes that are associated with each pool.

To display information on pools:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display pool information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Pools.
The Storage Pool Summary page is displayed.
4. Click a pool name for detailed information on that pool.
The Storage Pool Details page for the selected pool is displayed.
5. Go to Related Information and click any item for more information associated with the selected pool.
The Summary page for the selected item is displayed.

Creating a Storage Pool

To create a storage pool:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to create a storage pool.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Pools.
The Pool Summary page is displayed.
4. Click New.
The Create New Storage Pool page is displayed.
5. Specify a name for the new pool, using a maximum of 30 characters.
6. Enter a description of the new storage pool.
7. Select a storage profile for this pool, and click OK.

Modifying a Storage Pool

To modify a storage pool:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to modify a storage pool.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Pools.
The Storage Pool Summary page is displayed.
4. Click the storage pool that you want to modify.
The Storage Pool Details page for that pool is displayed.
5. Make the appropriate modifications, and click OK.

Deleting a Storage Pool

Deleting a storage pool will also delete all volumes associated with the pool.

To delete a storage pool:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array from which you want to delete a storage pool.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Pools.
4. Click the check box for the storage pool that you want to delete.
This enables the Delete button.
5. Click Delete.
The selected storage pool is no longer displayed in the Storage Pools table.

Configuring Storage Profiles

This section describes storage profiles. It contains the following subtopics:

- “About Storage Profiles” on page 113
- “Reference: Standard Storage Profiles” on page 114
- “Displaying Profile Information” on page 115
- “Creating a Storage Profile” on page 115
- “Modifying a Storage Profile” on page 116
- “Copying a Storage Profile” on page 116
- “Deleting a Storage Profile” on page 117

About Storage Profiles

A storage profile is set of attributes that you apply to a storage pool to allocate storage, instead of having to set each attribute individually. The system has a predefined set of storage profiles. The array comes with a number of preconfigured storage profiles to meet different requirements. You can choose a profile suitable for the application that is using the storage, or you can create a custom profile.

The array has a default storage profile with RAID-5 storage characteristics that is suitable for many storage applications. The default pool uses the default profile.

Each storage profile has the settings described in TABLE 4-7.

TABLE 4-7 Storage Profile Settings

Parameter	Value or Variable Type	Description
Name	Up to 32 characters	Unique identifier for the storage profile.
RAID Level	0, 1, 3, 5	RAID level configured across all disks within a virtual disk.
Segment Size	8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, 512 KB	Segment size for this profile.
Readahead	Enabled or Disabled	Read-ahead mode of the array. Cache read-ahead enables the controller to copy additional data blocks into cache while the controller reads and copies host requested data blocks from disk into cache.
Disk Type	ANY, FC, or SATA	Disk type.

TABLE 4-7 Storage Profile Settings (*Continued*)

Parameter	Value or Variable Type	Description
No. of Disks	1 to 14, or variable	The number of disks to be grouped together in a virtual disk. For example, if you create a storage pool with a profile that has the number of disks parameter set to a number, all virtual disks that are part of that storage pool must have the same number of disks. If the number of disks parameter is set to the Variable value you are prompted for the number of disks when storage is added to the pool.

Reference: Standard Storage Profiles

The array includes the storage profiles described in TABLE 4-8, all of which can be used with any disk type.

TABLE 4-8 Standard Storage Profiles

Name	RAID Level	Segment Size	Read-Ahead Mode	Disk Drive Type	No. of Disks
Default	RAID-5	512 KB	Enabled	FC	Variable
High_Capacity_Computing	RAID-5	512 KB	Enabled	SATA	Variable
High_Performance_Computing	RAID-5	512 KB	Enabled	FC	Variable
Mail_Spooling	RAID-1	512 KB	Enabled	FC	Variable
NFS_Mirroring	RAID-1	512 KB	Enabled	FC	Variable
NFS_Striping	RAID-5	512 KB	Enabled	FC	Variable
Oracle_DSS	RAID-5	512 KB	Enabled	FC	Variable
Oracle_OLTP	RAID-5	512 KB	Enabled	FC	Variable
Oracle_OLTP_HA	RAID-1	512 KB	Enabled	FC	Variable
Random_1	RAID-1	512 KB	Enabled	FC	Variable
Sequential	RAID-5	512 KB	Enabled	FC	Variable
Sybase_DSS	RAID-5	512 KB	Enabled	FC	Variable
Sybase_OLTP	RAID-5	512 KB	Enabled	FC	Variable
Sybase_OLTP_HA	RAID-1	512 KB	Enabled	FC	Variable

Displaying Profile Information

You can display summary and detail information on existing storage profiles. You can also display summary information about pools and volumes associated with each storage profile.

To display information on storage profiles:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see profile information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
4. Click a profile name for detailed information on that profile.
The Storage Profile Details page for that profile are displayed.
5. Go to Related Information and click an item for more information associated with the selected profile.
The Summary page for the selected item is displayed.

Creating a Storage Profile

To create a storage profile:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to create a profile.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
4. Click New.
The New Storage Profile page is displayed.

5. Specify a name for the new profile, using a maximum of 32 characters.
6. Specify a description for the new profile, using a maximum of 256 characters.
7. Specify the remaining profile attributes.
8. Review the specified information for the new storage profile. If you are satisfied, click OK.

Modifying a Storage Profile

You cannot modify the default factory profiles. If a profile's state is In Use, then you can change the profile name and description only. A profile is in the In Use state when it is associated with a storage pool.

To modify a storage profile:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to modify a profile.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
4. Click the storage profile that you want to modify.
The Storage Profile Details page for the selected profile is displayed.
5. Make the appropriate modifications, and click OK.

Copying a Storage Profile

You can copy a profile to another array or other arrays.

To copy a storage profile:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array for which you want to copy a profile.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
4. Select the check box to the left of the storage profile that you want to copy.
The Copy button is enabled.
5. Click Copy.
The Copy Storage Profile page is displayed.
6. Select the array or arrays to which you want to copy the profile, and click OK.

Deleting a Storage Profile

You cannot delete the default factory profiles, nor can you delete a profile that is in the In Use state (associated with a storage pool).

To delete a storage profile:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to delete a profile.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Logical Storage > Profiles.
The Storage Profile Summary page is displayed.
4. Select the check box to the left of the storage profile that you want to delete.
This enables the Delete button.
5. Click Delete.
The selected storage profile is removed from the Storage Profiles table.

Configuring Initiators

This section describes initiators. It contains the following subsections:

- “About Initiators” on page 118
 - “Displaying Initiator Information” on page 118
 - “Creating an Initiator” on page 119
 - “Deleting an Initiator” on page 120
-

About Initiators

An initiator is the protocol-specific physical port that initiates the I/O exchanges with the array. 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 data host is connected to the array by two host bus adapters (HBAs), the array handles the host as two different initiators. FC array LUN masking and mapping uses initiator port identifiers to authenticate storage customers.

When a new initiator is connected into the storage environment, the array detects it, and the management software displays it on the Initiator Summary page. To configure the initiator, select it from the list, add a description to identify it, assigning a host type, and then associating the initiator with a host. When you associate the host or its host group with a volume, you form a storage domain.

The array firmware retains all created or discovered WWNs until they are manually deleted. If the WWNs no longer apply because a server is moved or rezoned, delete the initiators with the affected WWNs from the Initiator Summary page.

To see the current initiators, go to the Initiator Summary page, as described in “Displaying Initiator Information” on page 118.

Displaying Initiator Information

You can display summary and detail information about existing initiators. You can also display summary information about mapped hosts associated with each initiator.

To display information on initiators:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display initiator information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Initiators.
The Initiator Summary page is displayed.
4. Click an initiator name for detailed information on that initiator.
The Initiator Details page for the selected initiator is displayed.
5. Go to Related Information and click Mapped Volumes for more information associated with the selected initiator.
The Summary page for the selected item is displayed.

Creating an Initiator

When you create an initiator, you name it and associate it to one specific host, and set the host type.

You will need the World Wide Name (WWN) for the initiator that you want to associate with a volume.

To create an initiator:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to create an initiator.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Initiators.
The Initiator Summary page is displayed.
4. Click New.
The New Initiator page is displayed.
5. Specify a name for the new initiator, using a maximum of 30 characters.

6. Select an existing World Wide Name (WWN) or specify a new WWN for the initiator.

The list displays only existing WWNs that are currently unassigned. The WWN that you select is automatically displayed in the Type a New WWN field. If no unassigned WWNs exist, the list is empty and you must enter a new WWN in the Type a New WWN field.

7. Select the host with which you want to associate the initiator.
8. Select the host type, usually the data host operating system, of the initiator.
9. Review the specified information. If you are satisfied, click OK.

Deleting an Initiator

To delete an initiator:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to delete an initiator.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Initiators.
The Initiator Summary page is displayed.
4. Select the initiator that you want to delete.
This enables the Delete button.
5. Click Delete.
6. Confirm the deletion.
The initiator disappears from the Initiator Summary page.

Configuring Host Groups and Hosts

This section describes storage host groups and hosts. It contains the following subsections:

- “About Host Groups” on page 121
- “Managing Host Groups” on page 122
- “About Hosts” on page 127
- “Managing Hosts” on page 127

About Host Groups

A host group is collection of one or more data hosts. The hosts in a host group access a storage array volume when the host group and the specified volume are included in a storage domain, which is a logical entity used to partition storage.

When you create a host group, the management software automatically includes it in the default storage domain. The default storage domain contains all hosts and host groups without explicit mappings and enables them to share access to volumes assigned a default logical unit number (LUN) mapping during volume creation. When you explicitly map a host group to a specific volume and LUN, the host group and volume are removed from the default storage domain and create a separate storage domain, which enables only the hosts in the group to share access to the specified volume.

A host group can be mapped to up to 256 volumes or snapshots.

A host can be a member of a host group. As a member of the group, the host has access to all volumes that are mapped to the host group, in addition to any volumes that are mapped directly to the host. For example, suppose that Host 1 is mapped to volume1, volume2, and volume3 and Host Group A is mapped to volume4 and volume5. If we add Host 1 to Host Group A, Host 1 now has access to volume1, volume2, volume3, volume4, and volume5.

Managing Host Groups

You can create host groups, you can add hosts to and remove hosts from a host group, you can map a host group to a volume, and you can display information about existing host groups and the mappings associated with them.

Displaying Host Group Information

You can display summary and detail information on existing host groups. You can also display summary information about the hosts and volumes that are associated with each host group.

To display information on host groups:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display host group information.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Physical Storage > Host Groups.
The Host Group Summary page is displayed.
4. Click a host group name for detailed information on that host group.
The Host Group Details page for the selected host group is displayed.
5. Go to Related Information and click any item for more information associated with the selected host group.
The Summary page for the selected item is displayed.

Creating a Host Group

When you create a host group, it is automatically included in the default storage domain.

To create a new host group:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array on which you want to create a host group.

The navigation pane and the Volume Summary page for that array are displayed displayed.

3. In the navigation pane, choose Physical Storage > Host Groups.

The Host Group Summary page is displayed.

4. Click New.

The New Host Group page is displayed.

5. Specify a name for the new host name, using a maximum of 30 characters.

6. Specify which of the available hosts you want to add to the group by double-clicking the host name or names. Click Add All to add all of the hosts defined on the storage array to the host group.

Note: A host can be in at most one host group at a time. If a host is already included in one host group, adding that host to another host group will move the host from one group to the other.

7. Click OK.

The Host Group Summary page is displayed listing the new host group.

Adding/Removing Host Group Members

A host can be a member of at most one host group. To move a host from one host group to another, you must first remove it from its current host group before you can add it to another host group.

If a host is mapped to one or more volumes and you move the host from one host group to another, the mappings specific to the host are retained but mappings associated with the host group from which the host was moved are not retained. Instead, the host inherits any mappings associated with the new host group to which it is moved.

If you move a host from a host group that has no mappings and that has host members without specific mappings, the host group and hosts are moved to the default storage domain.

For example, suppose that two host groups are created on an array, Host Group 1 and Host Group 2. Host Group 1 has three member hosts: Host1, Host2, and Host3. Host Group 1 has no mappings, and Host1 is the only of its hosts with any mappings. Host Group 2 has two member hosts: Host4 and Host5. Host Group 2 is mapped to one volume, and Host4 and Host5 have no specific mappings.

When you move Host1 from Host Group 1 to Host Group 2, the following occurs:

- Host1 keeps its specific mappings and inherits the mapping associated with Host Group 2.
- Host Group 1, Host2, and Host3 become part of the default storage domain, because they have no specific mappings.
- Host4 and Host5 are unaffected because they have no specific mappings.

To add hosts to or remove hosts from a host group:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to add or remove host group members.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Host Groups.
The Host Group Summary page is displayed.
4. Click the name of the host group that you want to modify.
The Host Group Details page is displayed.
5. Click Add/Remove Members.
The Add/Remove Host Group Members page is displayed.
6. Specify which of the available hosts you want to add to or remove from the group. To move an individual host, double-click the host name. To add all of the available hosts to the host group, click Add All. To remove all of the selected hosts from the host group, click Remove All.
7. Click OK.
The Host Group Summary page is displayed with the updated information.

Deleting a Host Group

Deleting a host group will also delete all associated hosts and all associated initiators. In addition, all associated volume and snapshot mappings will be removed.

To delete a host group:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.

2. Click the array on which you want to delete a host group.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Host Groups.
4. Select the check box next to the name of the host group that you want to delete.
This enables the Delete button.
5. Click Delete.
The Host Group Summary page is displayed with the updated information.

Mapping a Host Group to a Volume

Before you can map a host group to one or more volumes, the host group must have at least one host associated to at least one initiator.

When you explicitly map a host group to a volume, the management software removes the host group and specified volume from the default storage domain and creates a separate storage domain.

To map a host group to one or more volumes:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to map a host group to a volume.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Host Groups.
The Host Group Summary page is displayed.
4. Click the name of the host group that you want to map.
The Host Group Details page is displayed.
5. Click Map.
The Map Host Group page is displayed.
Note: The Map button is disabled when 256 volumes are assigned to the selected host group.
6. Select the check box next to the names of one or more volumes that you want to map to this host group.

7. Specify the logical unit number (LUN) (0 to 255 for the Sun StorageTek 6130 and 6140 arrays; 0 to 31 for the Sun StorageTek 6540 array) for each volume that you want to map to this host group.

8. Click OK.

The management software maps the volume to the host group and displays the Host Group Details page with the updated information.

Unmapping a Host Group From a Volume

Unmapping a host group from a volume removes the mapping between the volume and the host group (including all host group members). The volume is moved to an unmapped state. If this mapping was the last mapping to the host group and there are no other volume mappings to the hosts in the host group, this host group and all its members become part of the default storage domain.

To unmap a host group from a volume:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to unmap a host group from a volume.

The navigation pane and the Volume Summary page for that array are displayed displayed.

3. In the navigation pane, choose Physical Storage > Host Groups.

The Host Group Summary page is displayed.

4. Click the name of the host group that you want to unmap.

The Host Group Details page is displayed.

5. Click Unmap.

The Mapped Volumes and Snapshots page is displayed.

6. Select the check box next to volume or volumes that you want to unmap from this host group.

7. Click OK.

About Hosts

A host represents a data host, which is a computer that is attached to the storage array and generates data to be saved in a storage array. The data host sends data to the storage array using a Fibre Channel (FC) connection. A host must have one or more associated initiators to enable you to map volumes to the host.

A host can access a storage array volume when both the host and the volume are included in a storage domain, which is a logical entity used to partition storage.

When you create a host, the management software automatically includes it in the default storage domain. The default storage domain contains all hosts and host groups without explicit mappings and enables them to share access to volumes assigned a default logical unit number (LUN) mapping during volume creation. When you explicitly map a host to a specific volume and LUN, the host and volume are removed from the default storage domain and create a separate storage domain, which enables only that host to access the specified volume.

An individual host can be explicitly mapped to one volume. A host can also be a member of a host group. As a member of the group it has access to all volumes that are mapped to the hostgroup, in addition to any volumes that are mapped directly to the host. For example, Host 1 is mapped to volume1, volume2, and volume3 and Hostgroup A is mapped to volume4 and volume5. If we add Host 1 to Hostgroup A, Host 1 now has access to volume1, volume2, volume3, volume4, and volume5.

Managing Hosts

You can create hosts, you can map a host to a volume, and you can display information about existing hosts and the mappings associated with them.

Displaying Host Information

You can display summary and detail information on existing hosts. You can also display summary information about the initiators and volumes that are associated with each host.

To display information on hosts:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to display host information.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Hosts.
The Host Summary page is displayed.
4. Click a host name for detailed information on that host.
The Host Details page for the selected host is displayed.
5. Go to Related Information and click any item for more information associated with the selected host.
The Summary page for the selected item is displayed.

Creating a Host

When you create a host, you are associating a host name with which to identify a data host to the array; you are not creating a host on the network. By default, the new host is included in the default storage domain.

To create a new host:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to create a host.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Hosts.
The Hosts Summary page is displayed.
4. Click New.
The Create New Host page is displayed.
5. Specify a name for the new host, using a maximum of 30 characters.
6. Optionally select a host group for the new host if you want the new host to share access to the volume with the other hosts in the group.

7. Click OK.

The host is created and the Host Summary page is displayed, listing the new host.

Deleting a Host

Deleting a host will delete all associated initiators and remove all mappings to volumes and snapshots associated with the host.

To delete a host:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to delete a host.

The navigation pane and the Volume Summary page for that array are displayed.

3. In the navigation pane, choose Physical Storage > Hosts.

4. Select the check box next to name of the host that you want to delete, and the Delete button becomes active.

5. Click Delete.

Mapping a Host to a Volume

Up to 256 volumes can be explicitly mapped to an individual host. You can also include the same host in one or more host groups which are mapped to other volumes.

Before you map a host to a volume, make sure that it is associated with an initiator.

To map a host to a volume:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to map a host to a volume.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Physical Storage > Host.

The Host Summary page is displayed.

4. Click the name of the host that you want to map.
The Host Details page is displayed.
5. Click Map.
The Map Host page is displayed.
Note: The Map button is disabled when maximum number of volumes are assigned to the selected host.
6. Click the name of the volume that you want to map to the host.
7. Specify the logical unit number (LUN) (0 to 255 for Sun StorageTek 6140 and 6130 arrays, and 0 to 31 for the Sun StorageTek 6540 array) of this volume to map to the host.
8. Click OK.
A confirmation message is displayed.

Unmapping a Host From a Volume

Unmapping a host from a volume removes the mapping between the host and the volume. The volume is moved to the unmapped state.

If this mapping was the last mapping to the host and the host is not a member of a host group, this host becomes part of the default storage domain.

If this mapping was the last mapping to the host and the host is member of a host group and there are no other explicit mappings to the host group or any other host in the host group, then this host group and all of its members become part of the default storage domain.

To unmap a host from a volume:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to unmap a host from a volume.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Host.
The Host Summary page is displayed.
4. Click the name of the host that you want to unmap.
The Host Details page is displayed.

5. Click Unmap.

The Mapped Volumes and Snapshots page is displayed.

6. Click the volume that you want to unmap from this host.

7. Click OK.

Managing Trays and Disk Drives

This section describes storage trays and disk drives. It contains the following subsections:

- “About Trays and Disk Drives” on page 132
- “Displaying General Controller Information” on page 133
- “Displaying Controller Performance Statistics” on page 133
- “Testing Controller Communication” on page 134
- “Displaying Port Information” on page 134
- “Displaying Tray Information” on page 135
- “Displaying Disk Information” on page 136
- “Managing Hot-Spare Drives” on page 136

About Trays and Disk Drives

Storage trays are identified according to whether they contain a redundant array of independent disks (RAID) controller:

- A controller tray, available in the Sun StorageTek 6140 and 6130 arrays, is a storage tray with hardware RAID management built into the controller. For controller information about the Sun StorageTek 6540, see the *Sun StorageTek 6540 Release Notes*.
- An expansion tray is a storage tray that has disks only and no controller. This type of tray is managed by a controller tray.

In the Sun StorageTek 6140 and 6130 arrays, each tray has a minimum of 5 and a maximum of 14 disk drives, numbered 1 to 14, from left to right. With the maximum of 8 trays in a cabinet, you can have up to 112 drives. When a drive is installed, the drive and tray slot designations are set automatically.

For the Sun StorageTek 6140 array, you can use either 4-gigabyte or 2-gigabyte Fibre Channel (FC) or Serial Advanced Technology Attachment (SATA) disk drives. Although you can mix drive types in the same tray, for optimal performance Sun recommends that all disk drives in a tray must be of the same type.

For the Sun StorageTek 6130 and 6540 arrays, you can use 2 gigabyte Fibre Channel (FC) or Serial Advanced Technology Attachment (SATA) disk drives. All disk drives in a tray must be of the same type.

Other characteristics of the disk drives are described in TABLE 4-9.

TABLE 4-9 Disk Drive Characteristics

Feature	Description
Drive capacity	72 GB/15 K, 72 GB/10 K, 146 GB/10 K
Mixed drive sizes	Allowed but not recommended. If drives of different sizes are mixed in a storage tray, all the drives used in a volume are treated as the smallest drive. For example, in a volume with a drive of 36 gigabytes and a drive of 146 gigabytes, only 36 gigabytes of the second drive is used. When the system is operating with mixed drive sizes, the hot-spare must be the size of the largest drive in the tray.
Array hot-spare	Available as a spare to any virtual disk in any tray in the array configuration.

Displaying General Controller Information

To display general information on controllers:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click an array.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Controllers.
The Controller Summary page is displayed.

Displaying Controller Performance Statistics

To display information about the controllers' performance:

1. Click Sun StorageTek Configuration Service.

- The Array Summary page is displayed.
2. Click the array for which you want to see performance statistics.
The navigation pane and the Volume Summary page for that array are displayed.
 3. In the navigation pane, choose Physical Storage > Controllers.
The Controller Summary page is displayed.
 4. Click View Performance Statistics.
The Performance Statistics Summary - Controllers page is displayed.
-

Testing Controller Communication

You can determine whether the management host has direct Ethernet connectivity to a specific controller on an array.

To test the communication between the management host and the array controller:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array on which you want to test the communication between the controller and the management host.
The navigation pane and the Volume Summary page for that array are displayed displayed.
3. In the navigation pane, choose Physical Storage > Controllers.
The Controller Summary page is displayed.
4. Click Test Communications.
An informational message is displayed that identifies the controller, the Ethernet port, the IP address, and whether the array has passed or failed the communication test.

Displaying Port Information

You can display summary and detail information about the Fibre Channel (FC) ports.

To display information on FC ports:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display port information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Ports.
The Port Summary page is displayed.
4. Click a port name for detailed information on that port.
The Port Details page for the selected port is displayed.

Displaying Tray Information

You can display summary and detail information about existing storage trays. You can also display summary information about the disks that are associated with each tray.

To display information on trays:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display tray information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Trays.
The Tray Summary page is displayed.
4. Click a tray name for detailed information on that tray.
The Tray Details page for the selected tray is displayed.
5. Go to Related Information and click any item for additional information associated with the selected tray.
The Summary page for the selected item is displayed.

Displaying Disk Information

You can display summary and detail information on existing disks.

To display information on disks:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to display disk information.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Disks.
The Disk Summary page is displayed.
4. Click a disk name for detailed information on that disk.
The Disk Details page for the selected disk is displayed.
5. Go to Related Information and click any item for more information associated with the selected disk.
The Summary page for the selected item is displayed.

Managing Hot-Spare Drives

This section describes how to manage hot-spare drives. It includes the following subsections:

- “About Hot-Spare Drives” on page 136
- “Planning Hot-Spares” on page 137
- “Assigning and Unassigning Hot-Spares Manually” on page 137
- “Assigning and Unassigning Hot-Spares Automatically” on page 138

About Hot-Spare Drives

A hot-spare is a disk drive, containing no data, that acts as a standby in the storage array in case a drive fails in a RAID-1, RAID-3, or RAID-5 volume. The hot-spare adds another level of redundancy to the storage array. The array supports up to 15 hot-spares.

If a hot-spare is available when a disk drive fails, the hot-spare is automatically substituted for the failed disk drive, without intervention. The controller uses redundancy data to reconstruct the data from the failed drive onto the hot-spare. When you have physically replaced the failed disk drive, the data from the hot-spare is copied back to the replacement drive. This is called copyback.

If you do not have a hot-spare, you can still replace a failed disk drive while the storage array is operating. If the disk drive is part of a RAID-1, RAID-3, or RAID-5 volume group, the controller uses redundancy data to automatically reconstruct the data onto the replacement disk drive. This is called reconstruction.

Planning Hot-Spares

A disk drive should be assigned as a hot-spare only if it meets the following criteria:

- The hot-spare must have a capacity that is equal to or greater than the capacity of the largest drive on the storage array. If a hot-spare is smaller than a failed physical disk drive, the hot-spare cannot be used to rebuild the data from the failed physical disk drive.
- The hot-spare drive must be the same type of physical disk as the physical disk drive that failed. For example, a SATA hot-spare cannot replace a Fibre Channel physical disk drive.
- The disk drive's role must be unassigned, its state must be enabled, and it must have an optimal status.

Information about individual disk drives is available on the Disk Summary page.

Assigning and Unassigning Hot-Spares Manually

To manually assign or unassign a hot-spare:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to assign or unassign a hot-spare.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Physical Storage > Disks.
The Disk Summary page is displayed.
4. Select the disk drive you want to assign or unassign as a hot-spare drive.

5. Do one of the following:
 - To assign the selected disk as a hot-spare, click Assign Hot-Spare.
 - To unassign the selected disk as a hot-spare, click Unassign Hot-Spare.The Disk Summary page is displayed with the updated information.

Assigning and Unassigning Hot-Spares Automatically

When the management software assigns the hot-spares, it balances the request for spares across all trays within the array and ensures that the hot-spare drive is of the same type as the other disks in the same tray. It also verifies that the disk drives are unassigned, enabled, and in the optimal state.

To enable the management software to assign or unassign a hot-spare:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to assign a hot-spare.
The navigation pane and the Volume Summary page are displayed.
3. In the navigation pane, choose Administration > General Setup.
The General Setup page is displayed.
4. In the Array Hot Spares Change field, select the number of hot-spares that you want to assign to this array.
5. Click OK.
The management software assigns or unassigns the specified number of hot-spares, balancing the selection among trays within the array.

Configuring Storage Domains

This section describes storage domains. It contains the following subsections:

- “About Storage Domains” on page 139
- “Preparing to Create a Storage Domain” on page 140
- “Determining How to Create a Storage Domain” on page 140
- “Enabling the Storage Domain Feature” on page 141
- “About Performing Mapping Functions” on page 142

About Storage Domains

A storage domain, also called a set or a storage partition, is a logical entity used to partition storage. To create a storage domain after volume creation, you must define a single host or a collection of hosts (called a host group) that will access the storage array. Then, you will need to define a volume-to-logical unit number (LUN) mapping, which will allow you to specify the host or host group that will have access to a particular volume in your storage array. The storage domain designates that only the selected host or host group has access to that particular volume through the assigned LUN.

When the storage domain consists of a volume mapped to a host group, it can enable hosts with different operating systems (heterogeneous hosts), to share access to a storage volume. A host within a host group can be mapped separately to a different volume.

A storage domain can contain up to 256 volumes. A volume can be included in only one storage domain and each LUN, from 0 to 255, can only be used once per storage domain.

Note: Not all operating systems support up to 256 LUN IDs. See the documentation for your operating system for more information.

A default storage domain exists to include the following:

- All host groups and hosts that are not explicitly mapped to a volume.
- All volumes that have a default volume-to-LUN mapping assigned.
- All automatically detected initiators.

Any volumes within the default storage domain can be accessed by all hosts and host groups within that storage domain.

Creating an explicit volume-to-LUN mapping for any host or host group and volume within the default storage domain causes the management software to remove the specified host or host group and volume from the default storage domain and create a new separate storage domain.

The array management software supplies eight storage domains. You must obtain a license for additional storage domain support.

Preparing to Create a Storage Domain

You create a storage domain by explicitly mapping a volume or snapshot to a host or host group. Before creating a storage domain, you should perform the following tasks:

1. Create one or more hosts.
2. Create one or more initiators for each host.

A host included in a storage domain, either as an individual host or as a member of a host group, can have access to storage volumes only if it is associated with one or more initiators.

3. Create a host group, if needed, and add one or more hosts to it.
4. Create one or more volumes.

Determining How to Create a Storage Domain

There are several ways to create a storage domain:

- **During standard volume or snapshot creation**

When you create a volume or snapshot, the wizard prompts you to map the volume either now or later.

Create a storage domain at this time by mapping the volume to either the default storage domain or to hosts or hosts that you have already created.

- **After the total storage array capacity has been configured into volumes**

- From the Mapping Summary page, launch the Create New Mappings wizard.

Create a storage domain using the wizard if you want to map an existing volume or snapshot, assign a logical unit number (LUN) to it, and then select an existing host or host group to which to map the volume or snapshot.

- From the Volume Summary or Volume Details page, you can click Map to map the volume to an existing host or host group.

Create a storage domain from the Volume Summary page if you want to map existing standard volumes to existing hosts or host groups. Use the Volume Details page if you want to map or remap a volume after reviewing the volume details.

- From the Host Details page or Host Group Details page, click Map to map a host or host group to an existing volume.

Create a storage domain from the Host Details page or Host Groups Details page if you want to map several standard volumes or snapshot volumes to one host or host group.

- From the Snapshot Summary page or Snapshot Details page, click Map to map a snapshot to an existing host or host group.

Create a storage domain from the Snapshot Summary page or Snapshot Details page if you want to map existing snapshot volumes to existing hosts or host groups. Use the Snapshot Details page if you want to map or remap the snapshot volume after reviewing the snapshot volume details.

Enabling the Storage Domain Feature

The Sun StorageTek Storage Domain is a premium feature. Licenses are available for 8, 16, or 64 storage domains. You can determine the number of configurable storage domains by obtaining the appropriate license.

To enable storage domains:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to enable additional storage domains.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Licensing.
The Licensable Feature Summary page is displayed.
4. Click Add License.
The Add License page is displayed.

5. Select the type of storage domain license that you have purchased from the License Type menu.

6. Enter the version number and the key digest, and click OK.

A message displays, indicating that the operation was successful and the Licensable Feature Summary page is displayed with the updated information.

About Performing Mapping Functions

The management software enables you to manage mappings from the Mapping Summary page. However, it also provides several other locations from which you can perform specific mapping functions or view mappings.

TABLE 4-10 shows the locations from which you can perform mapping operations.

TABLE 4-10 Mapping Matrix

Object to Be Mapped/Unmapped	Source of Mapping or Unmapping	Mapping Location
Volume	Host or host group	Map Volume page Mapping Summary page
Snapshot	Host or host group	Map Snapshot page Mapping Summary page
Host	Volume or snapshot	Map Host page Mapping Summary page
Host group	Volume or snapshot	Map Host Group page Mapping Summary page

Displaying All Mappings on the Array

To display all mappings on the array:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array for which you want to display all mappings.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Mappings.

The Mapping Summary page lists all mappings on the array.

Creating a Mapping

The management software enables you to map a volume or snapshot to a host or host group from the Mapping Summary page. In addition, you can perform more specific mappings from the following locations:

- Volume Summary page
- Snapshot Summary page
- Host Group Summary page
- Host Summary page

To map a volume or snapshot to a host or host group:

1. Click Sun StorageTek 6140 Configuration Service.

The Array Summary page is displayed.

2. Click the array for which you want to display all mappings.

The navigation pane and the Volume Summary page are displayed.

3. In the navigation pane, choose Mappings.

The Mappings Summary page lists all mappings on the array.

4. Click Map.

The Create New Mappings wizard is displayed.

5. Follow the steps in the wizard.

Click the Help tab in the wizard for more information.

Troubleshooting the System

This chapter describes troubleshooting. It contains the following subsection:

- “Basic Troubleshooting” on page 146

Basic Troubleshooting

This section described updating array components and troubleshooting the array:

- “About Patches, Firmware, and Software” on page 146
- “Backing Out Changes” on page 147
- “About Field-Replaceable Units” on page 147
- “About LED Status Indicators” on page 147
- “Checking LEDs” on page 156
- “Responding to an Array Boot-Up Failure” on page 157
- “Viewing Log Files” on page 157
- “Resetting a Storage Array Configuration” on page 157
- “Redistributing Volumes” on page 158
- “Resetting a Controller” on page 159
- “Reviving Virtual Disks and Disk Drives” on page 160

About Patches, Firmware, and Software

You can use the Sun Services web site or compact discs (CDs) to keep the array patches, software, and firmware up to date. Use these tools to do the following:

- Make sure all patches, firmware, and software are up to date at array startup.
- Establish a site process for making sure patches, firmware, and software are kept current.
- Monitor the levels manually at any time.

Check the release notes for each software release to check patch, firmware, and software changes and upgrades. Sun Services can make updates available from the web site or on CD.

The upgradable components are:

- Array firmware
- Storage Automated Diagnostic Environment software
- Controller image
- Disk image

Because components in the array 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 array operates properly. Then contact Sun Services with a description of the component's failed upgrade process.

Backing Out Changes

You can track and restore the following changes:

- Operating system upgrades
- Patch upgrades
- Firmware upgrades

Refer to the release notes for more information about this feature.

About Field-Replaceable Units

For the array, field-replaceable units (FRUs) can be replaced by Sun field engineers or by Sun-trained customer administrators.

To see a list of the hardware components that can be replaced at the customer site:

1. Go to the Java Web Console page and click Sun Storage Automated Diagnostic Environment.
2. In the top right of the page, click Service Advisor.

The Storage Automated Diagnostic Environment software displays the Service Advisor page. The Service Advisor combines hardware procedures with reports and automated steps using Storage Automated Diagnostic Environment features. Choose one of the types of FRUs or other options to see the procedures.

About LED Status Indicators

The following section describes the light-emitting diodes (LEDs) on the redundant array of independent disks (RAID) controller and the controller and expansion trays, including the following:

FIGURE 5-1 shows the LEDs and components on the front of the controller tray.

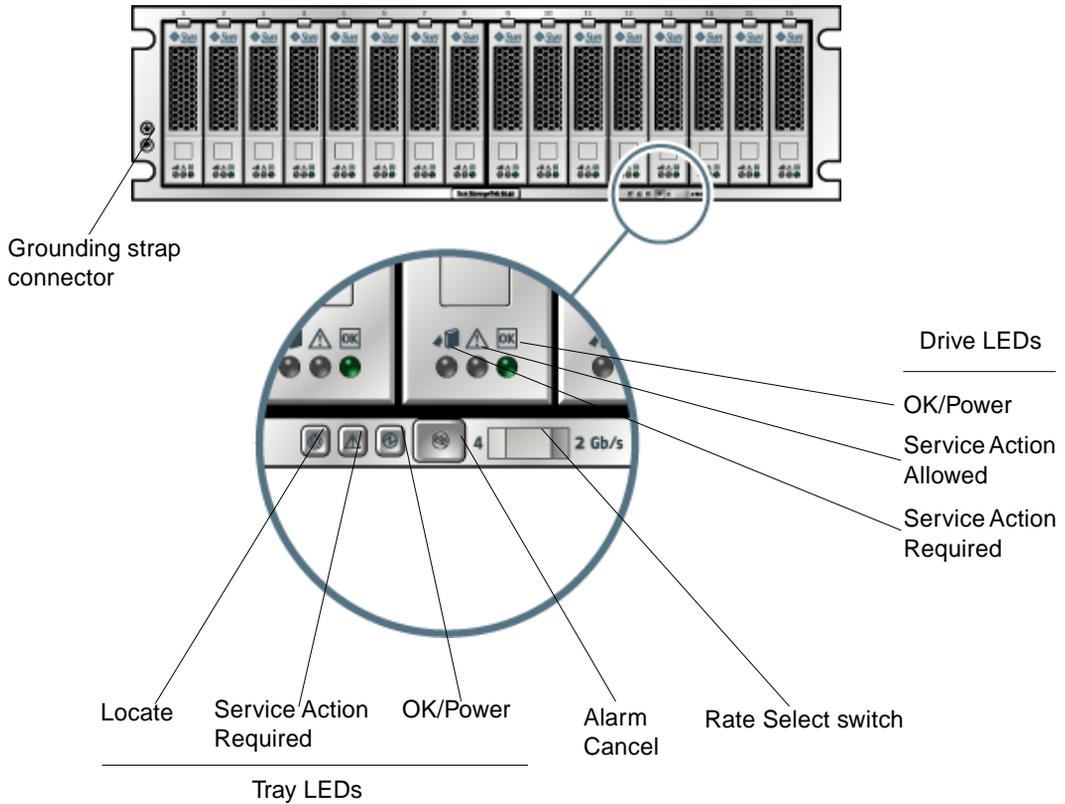


FIGURE 5-1 Controller Tray (Front View)

TABLE 5-1 describes the LEDs and components on the front of the controller tray.

TABLE 5-1 Tray LED Components (Front)

LED/Component	Description
<i>Drive LEDs</i>	
Service Action Allowed	Steady blue indicates that service action can be taken on the drive without adverse consequences. Off indicates that the drive is engaged and service cannot be implemented.
	

TABLE 5-1 Tray LED Components (Front) (*Continued*)

LED/Component	Description
Service Action Required	Steady amber indicates that the drive requires service. Off indicates that the drive does not require service.
OK	Steady green indicates that power is applied to the drive and the drive is functioning normally. Off indicates that power is not applied to the drive.
OK	Flash indicates that normal activity is in progress.
<i>Tray LEDs</i>	
Locate	Steady white identifies the tray after initiation from the management station.
Service Action Required	Steady amber indicates that the tray requires service. Off indicates that the tray does not require service.
OK/Power	Steady green indicates that power is applied to the tray and the tray is functioning normally. Off indicates that power is not applied to the tray.
	Flash indicates that normal activity is in progress.
<i>Tray Components</i>	
Alarm Cancel button	Reserved for canceling audio alarm functionality that is not currently activated. Use the management software to review alarms and events.
Rate Select switch	When the switch is in the left position, the link rate for the tray is 4 Gbits/second; When the switch is in the right position, the link rate for the tray is 2 Gbits/second. Set the switch to match the link rate of the drive with the lowest link rate in the tray.

TABLE 5-1 Tray LED Components (Front) (Continued)

LED/Component	Description
Grounding strap connector	Use this connector to connect a grounding strap to the tray before handling the tray or its components.

FIGURE 5-2 shows the LEDs and indicators at the back of the controller tray.

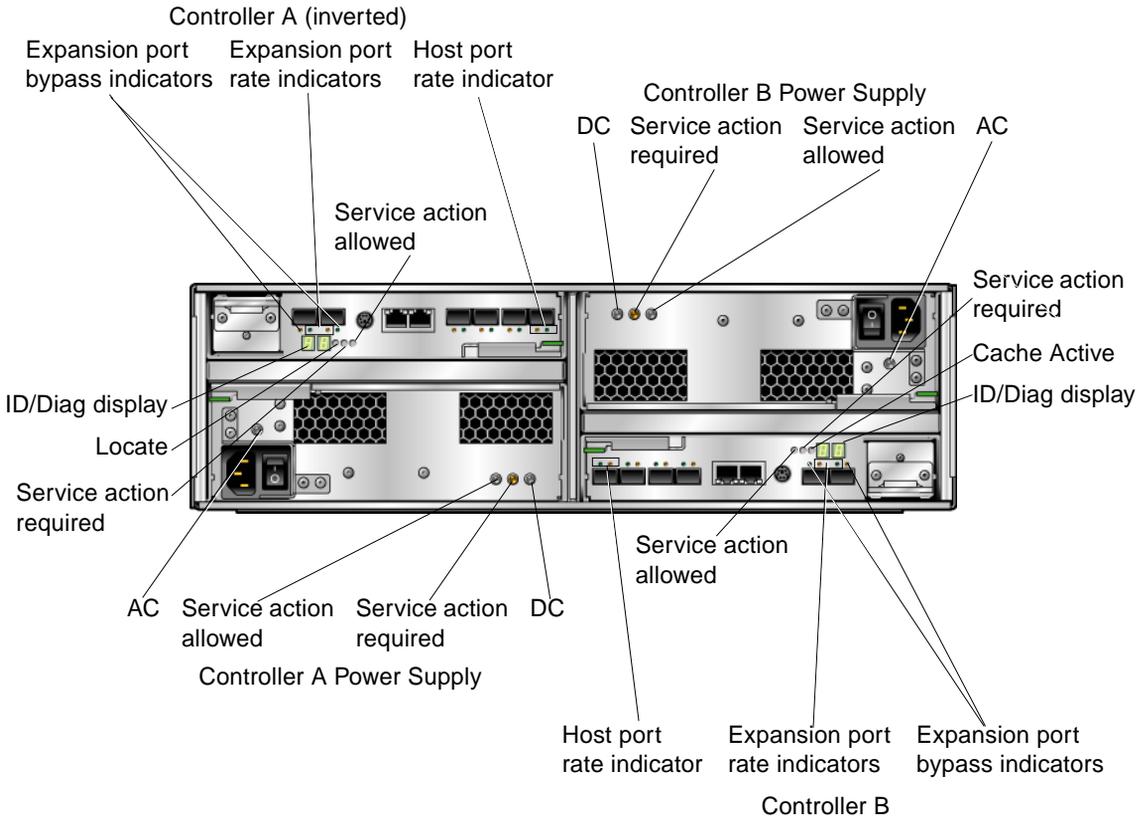


FIGURE 5-2 Controller Tray LEDs and Indicators (Back View)

TABLE 5-2 describes the LEDs and indicators on the back of the controller tray.

TABLE 5-2 Controller Tray LEDs and Indicators (Back)

LED/Indicator	Description
<i>Power Supply LEDs</i>	
DC 	On indicates that the correct DC power is being output from the controller power supply.
Service Action Required 	Steady amber indicates that the power supply requires service. Off indicates that the power supply does not require service.
Service Action Allowed 	Steady blue indicates that service action can be taken on the power supply without adverse consequences. Off indicates that the power supply is engaged and service action should not be implemented.
AC 	On indicates that AC power is being supplied to the controller power supply.
<i>Controller LEDs</i>	
ID/Diag display	Seven-segment readouts indicate the ID of the tray.
Cache Active 	Steady green indicates that data is in the cache. Off indicates that all data has been written to disk and the cache is empty.
Service Action Required 	Steady amber indicates that the controller requires service. Off indicates that the controller does not require service.
Service Action Allowed 	Steady blue indicates that service action can be taken on the controller without adverse consequences. Off indicates that the controller is engaged and service action should not be implemented.

TABLE 5-2 Controller Tray LEDs and Indicators (Back) (*Continued*)

LED/Indicator	Description
<i>Controller Indicators</i>	
Host Port Rate 	The combined display indicates the host port link rate for the tray: <ul style="list-style-type: none">• LED 1 On, LED 2 On – 4 Gbits/second• LED 1 Off, LED 2 On – 2 Gbits/second• LED 1 On, LED 2 Off – 1 Gbits/second (Not supported)
Expansion Port Rate 	The combined display indicates the expansion port link rate for the tray: <ul style="list-style-type: none">• LED 4 On, LED 2 Off – 4 Gbits/second• LED 4 Off, LED 2 On – 2 Gbits/second
Expansion Port Bypass 	Steady amber indicates that no valid device is detected and that the drive port is bypassed. Off indicates that there is no small form factor plug-in (SFP) transceiver installed or that the port is enabled.

FIGURE 5-3 shows the LEDs on the back of the expansion tray.

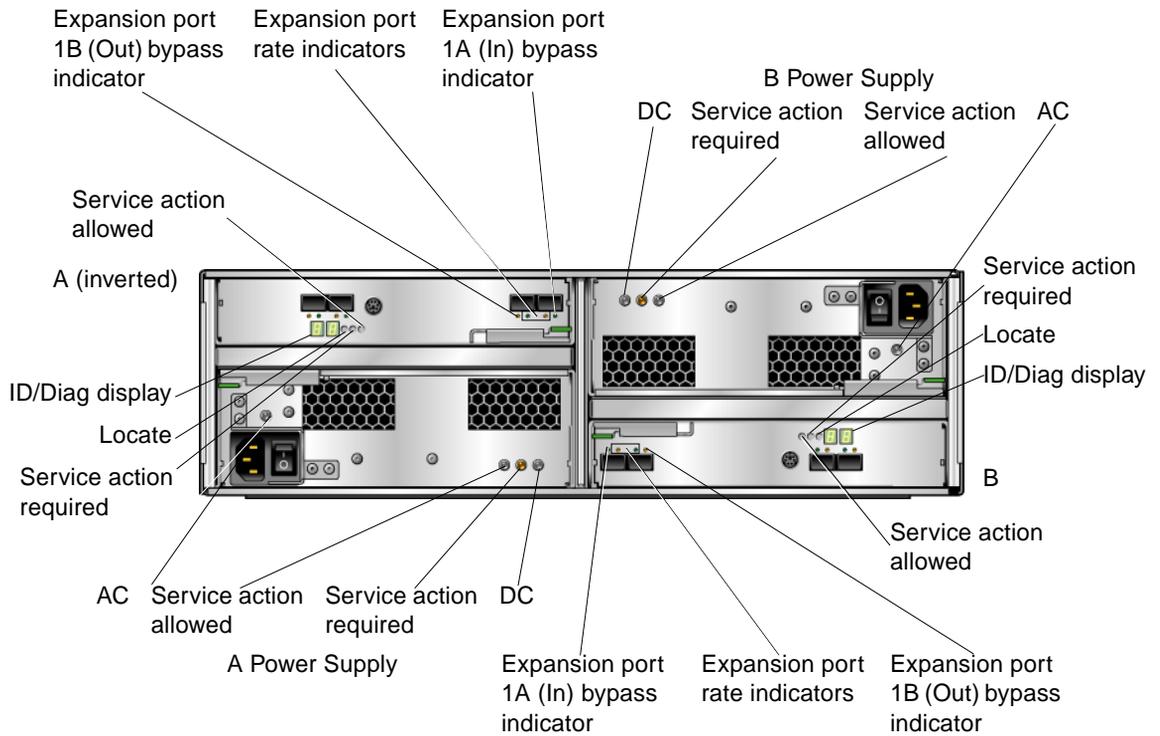


FIGURE 5-3 Expansion Tray LEDs and Indicators (Back)

TABLE 5-3 describes the LEDs and indicators on the back of the expansion tray.

TABLE 5-3 Expansion Tray LEDs and Indicators (Back)

LED/Indicator	Description
<i>Power Supply LEDs</i>	
DC	On indicates that the correct DC power is being output from the controller power supply.
	
Service Action Required	Steady amber indicates that the power supply requires service. Off indicates that the power supply does not require service.
	

TABLE 5-3 Expansion Tray LEDs and Indicators (Back) (*Continued*)

LED/Indicator	Description
<p>Service Action Allowed</p> 	<p>Steady blue indicates that service action can be taken on the power supply without adverse consequences. Off indicates that the power supply is engaged and service action should not be implemented.</p>
<p>AC</p> 	<p>On indicates that AC power is being supplied to the controller power supply.</p>
<p><i>Expansion Tray LEDs</i></p>	
<p>ID/Diag display</p>	<p>Seven-segment readouts indicate the ID of the tray.</p>
<p>Locate</p> 	<p>Steady white identifies the controller after initiation from the management station.</p>
<p>Service Action Required</p> 	<p>Steady amber indicates that the controller requires service. Off indicates that the controller does not require service.</p>
<p>Service Action Allowed</p> 	<p>Steady blue indicates that service action can be taken on the controller without adverse consequences. Off indicates that the controller is engaged and service action should not be implemented.</p>
<p><i>Expansion Tray Indicators</i></p>	
<p>Expansion Port Rate</p> 	<p>The combined display indicates the expansion port link rate for the tray:</p> <ul style="list-style-type: none"> • LED 4 On, LED 2 Off – 4 Gbits/second • LED 4 Off, LED 2 On – 2 Gbits/second
<p>Expansion Port Bypass</p> 	<p>Steady amber indicates that no valid device is detected and that the drive port is bypassed. Off indicates that there is no SFP installed or that the port is enabled.</p>

FIGURE 5-4 shows the location of the battery compartment on the controller and expansion trays and identifies the LEDs on the compartment.

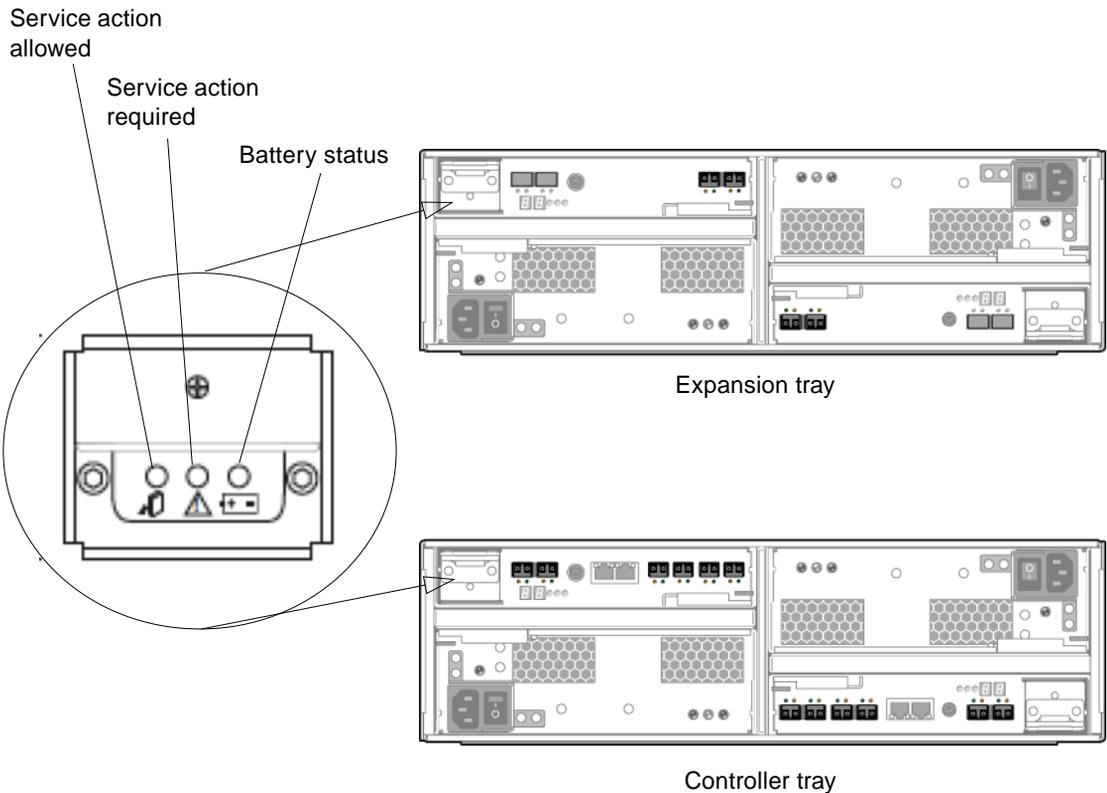


FIGURE 5-4 Battery Compartment LEDs

TABLE 5-4 describes the LEDs and indicators on the back of the expansion tray.

TABLE 5-4 Battery Compartment LEDs

LED/Indicator	Description
Service Action Allowed 	Steady blue indicates that service action can be taken on the power supply without adverse consequences. Off indicates that the power supply is engaged and service action should not be implemented.

TABLE 5-4 Battery Compartment LEDs (*Continued*)

LED/Indicator	Description
Service Action Required 	Steady amber indicates that the power supply requires service. Off indicates that service is not required for the power supply.
Battery Status 	Steady green on indicates that the battery is fully charged. Slow blink indicates that the battery is charging. Off indicates that the battery is discharged or off.

Checking LEDs

Light-emitting diodes (LEDs) provide status information on each array tray and its components. Green lights indicate a normal operating status, and amber lights indicate a hardware fault. Always check the status lights on the front and back of each array tray whenever you turn on the power. During power-up, the lights blink intermittently as the array and components complete the power-up process.

Caution: Electrostatic discharge can damage sensitive components. Touching the command tray or its components without using a proper ground may damage the equipment. To avoid damage, use proper antistatic protection before handling any components.

Use the following procedure to check the lights upon startup.

1. Before checking for faults, wait until the lights stop blinking. The Heartbeat light on the back of controller will continue to blink.
2. Check the Link Rate For Each Port
Link rate indicators for host and expansion ports are located on the back of the controller tray. Link rate indicators for expansion ports are located on the back of the expansion tray.
3. Check that no amber LEDs are lit on the front or back of the array. If an amber LED is lit, contact your Sun service representative.

Responding to an Array Boot-Up Failure

If the array does not boot after a shutdown, the failure could have several causes, both hardware and software. Check the messages reported on the console and in the event log to diagnose the problem.

Viewing Log Files

The array logs its events to the same log file maintained by the Sun Web Console.

To view system messages:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array for which you want to see messages.
The navigation pane and the Volume Summary page for that array are displayed.
3. In the navigation pane, choose Administration > Notification Management.
4. Click the link to browse the log files.

Resetting a Storage Array Configuration

Resetting a storage array configuration deletes the entire array configuration. In addition, it removes all data from the array and refreshes the settings to factory defaults. Reset the array configuration only to correct errors that cannot be corrected by any other method.

Resetting an array configuration produces the following results:

- Removes all volumes and virtual disks
- Clears the user-supplied array name
- Clears the array password
- Removes all data from the array

Note: Before you begin, be sure that all data on the storage array has been backed up to tape or another storage array.

To reset the array configuration:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array whose configuration you want to reset.

The navigation pane and the Volume Summary page for the selected array are displayed.

3. In the navigation pane, choose Administration > General Settings.

The General Setup page is displayed.

4. Click Reset Configuration.

A confirmation dialog box prompts you to confirm the selected action.

5. Click OK to continue.

Note: The array will not be available and will be in an error state until the reconfiguration process has been completed.

Redistributing Volumes

When volumes are created, they are assigned a preferred owner controller. The owner controller of the volume sometimes shifts from the preferred controller when the controller is being replaced or is undergoing a firmware download. In addition, multipath drivers move volumes from their preferred controller owner when a problem occurs along the data path between the host and the storage array. When you redistribute the volumes, you cause the volumes to move back to their preferred controller owners.

Note: Redistributing the volumes while an application is using the affected volumes causes I/O errors unless a multipath driver is installed on the data host. Therefore, before you redistribute volumes, verify either that the volumes are not in use or that there is a multipath driver installed on all hosts using the affected volumes.

To redistribute volumes:

1. Click Sun StorageTek Configuration Service.

The Array Summary page is displayed.

2. Click the array on which you want to redistribute the volumes.
The navigation pane and the Volume Summary page for the selected array are displayed.
 3. In the navigation pane, choose Administration > General Settings.
The General Setup page is displayed.
 4. Click Redistribute Volumes.
The Redistribute Volumes confirmation dialog box is displayed.
 5. Click OK.
-

Resetting a Controller

The ability to reset the controller is a recovery option available in the event of serious controller errors. Resetting a controller makes it unavailable for I/O until the reset has been completed. If a host is using volumes owned by the controller being reset, the I/O directed to the controller will be rejected. Before resetting the controller, either verify that the volumes owned by the controller are not in use or ensure that there is a multipath driver installed on all hosts using these volumes.

To reset a controller:

1. Click Sun StorageTek Configuration Service.
The Array Summary page is displayed.
2. Click the array whose controller you want to reset.
The navigation pane and the Volume Summary page for the selected array are displayed.
3. In the navigation pane, choose Physical Storage > Controllers.
The Controller Summary page is displayed.
4. Click Reset Controller for the controller you want to reset.
A confirmation dialog box is displayed.
5. Click OK.

Reviving Virtual Disks and Disk Drives

The Sun StorageTek Configuration Service enables you to more closely manage the states of the virtual disks and physical disk drives and recover from failures.

If one or more disk drives in a virtual disk are in the Failed state, you can recover from this failure by reviving the virtual disk. Reviving a virtual disk automatically revives the failed disk drives included in the virtual disk, assuming that the disk drives are still usable.

You should manually revive an individual disk drive only if reviving the virtual disk fails to revive any failed disk drives in the virtual disk.

Caution: These tasks must be performed in a specific order and only under the direct supervision of a Sun Customer and Technical Support representative.

From the Virtual Disk Details page, you can:

- Revive a virtual disk

If one or more disk drives in a virtual disk are in the failed state, you can recover from this failure by reviving the virtual disk. Reviving a virtual disk automatically revives the failed disk drives included in the virtual disk.

- Place a virtual disk offline

Placing a virtual disk offline disables the virtual disk.

- Place a virtual disk online

Placing a virtual disk online enables the virtual disk.

From the Disk Details page you can:

- Revive the disk drive

If reviving a virtual disk fails to revive an individual disk drive, you can manually revive that disk drive.

- Reconstruct the disk drive

You can reconstruct a disk drive only when it meets the following criteria:

- The disk drive is assigned to a RAID-1, -3, or -5 virtual disk.
- The disk drive has a Failed or Replaced status and was not automatically revive during a virtual disk revive operation.

- Fail the disk drive

Failing the disk drive disables the drive.

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.

A

- agent** The component of the system monitoring and diagnostic software that collects health and asset information about the array.
- alarm** A type of event that requires service action. See also [event](#).
- alert** A subtype of an event that requires user intervention. The term *actionable event* often describes an alert. 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.
- array hot-spare** A disk that serves as a hot spare within an array as part of a storage pool; a reserve disk that can be made available to all virtual disks within an array. See also [hot-spare](#).

B

- block** The amount of data sent or received by the host per I/O operation; the size of a data unit.

C

- capacity** The amount of storage you must allocate to storage elements, including volumes, pools, and virtual disks. Capacity planning should include allocations for volume snapshots and volume copies.
- CLI** Command-line interface. The SSCS command-line interface is available from the remote CLI client or through an SSCS directory on the Solaris Operating System management software station.

- controller tray** A tray with an installed redundant RAID controller pair. In a Sun StorageTek 6140 Array, 1x1, 1x2, 1x3, 1x4, 1x5, 1x6, and 1x7 array types are available.
- control path** The route used for communication of system management information, usually an out-of-band connection.
- customer LAN** See [site LAN](#).

D

- 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 array (direct attach storage, or DAS) or can be connected to an external switch that supports multiple data hosts (storage area network, or SAN). See also [host](#).
- data path** The route taken by a data packet between a data host and the storage device.
- direct attached storage (DAS)** A storage architecture in which one or two hosts that access data are connected physically to a storage array.
- disk** A physical drive component that stores data.

E

- event** A notification of something that happened on a device. There are many types of events, and each type describes a separate occurrence. See also [alarm](#) and [alert](#).
- expansion tray** A tray that does not have a RAID controller, used to expand the capacity of an array. This type of tray must be attached to a controller tray to function.
- extent** A set of contiguous blocks with consecutive logical addresses on a physical or virtual disk.

F

- failover and recovery** The process of changing the data path automatically to an alternate path.
- fault coverage** The percentage of faults detected against all possible faults or against all faults of a given type.
- FC** See [Fibre Channel \(FC\)](#).
- Fibre Channel (FC)** A set of standards for a serial I/O bus capable of transferring data between two ports at up to 100 megabytes/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.

Fibre Channel 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). The Fibre Channel switches are used to expand the number of servers that can connect to a particular storage port. 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.

FRU See [field-replaceable unit \(FRU\)](#).

H

HBA See [host bus adapter \(HBA\)](#).

host As a function of the Sun StorageTek 6140 Array configuration, a representation of a data host that is mapped to initiators and volumes to create a storage domain. See also [data host](#), [initiator](#).

**host bus adapter
(HBA)**

An I/O adapter that connects a host I/O bus to a computer's memory system. Abbreviated HBA. Host bus adapter is the preferred term in SCSI contexts. Adapter and NIC are the preferred terms in Fibre Channel contexts. The term NIC is used in networking contexts such as Ethernet and token ring. See also [initiator](#).

host group A group of hosts with common storage characteristics that can be mapped to volumes. See also [host](#).

hot-spare The drive used by a controller to replace a failed disk. See also [array hot-spare](#).

I

in-band traffic System management traffic that uses the data path between a host and a storage device. See also [out-of-band traffic](#).

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.

IOPS A measure of transaction speed, representing the number of input and output transactions per second.

L

LAN Local area network.

logical unit (LUN) The SCSI identifier for a volume as it is recognized by a particular host. The same volume can be represented by a different LUN to a different host.

LUN See [logical unit \(LUN\)](#).

M

MAC Address See [media access control \(MAC\) address](#).

management host A Solaris host serving the configuration, management, and monitoring software for the Sun StorageTek 6140 Array. The software on the station can be accessed with a browser interface or with a remote scripting command-line interface (CLI) client to access the SCS CLI commands.

master/alternate

master A design for reliability that uses redundant configuration. Array configurations share master/alternate master configurations: each array configuration has two controller trays that are grouped as one host. In each case, the master component uses the IP address and name. If the master fails, the alternate master assumes the IP address and name and takes over the master's functions.

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.

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.

multipathing A design for redundancy that provides at least two physical paths to a target.

O

out-of-band traffic System management traffic outside of the primary data path that uses an Ethernet network. See also [in-band traffic](#).

P

PDU See [power distribution unit \(PDU\)](#).

pool See [storage pool](#).

power distribution unit (PDU) The assembly that provides power management for the system. The redundant design uses two PDUs in each system so that the system's data path continues to function if one of the PDUs fails.

profile See [storage profile](#).

provisioning The process of allocation and assignment of storage to hosts.

R

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. Also, a phrase adopted from the 1988 SIGMOD paper A Case for Redundant Arrays of Inexpensive Disks.

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.

S

SAN See [storage area network \(SAN\)](#).

site LAN The local area network at your site. When the system is connected to your LAN, the system can be managed through a browser from any host on the LAN.

snapshot A copy of a volume's data at a specific point in time.

SSCS Sun Storage Command System, the command-line interface (CLI) that can be used to manage the array.

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 Automated Diagnostic Environment

An automated fault management system (FMS) for storage area network (SAN) devices. The FMS provides health and telemetry for Sun SAN devices.

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 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 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 tray** An enclosure containing disks. A tray with dual RAID controllers is called a controller tray; a tray without a controller is called an expansion tray.
- 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. 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).

T

- target** The system component that receives a SCSI I/O command. (SNIA).
- thin-scripting client** See [remote scripting CLI client](#).
- tray** See [storage tray](#).

V

- 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.
- 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.
- volume snapshot** See [snapshot](#).

W

WWN World Wide Name. A unique 64-bit number assigned by a recognized naming authority such as the Institute of Electrical and Electronics Engineers (IEEE) that identifies a connection (device) or a set of connections to the network. The World Wide Name (WWN) is constructed from the number that identifies the naming authority, the number that identifies the manufacturer, and a unique number for the specific connection.

Index

A

- account passwords, changing 34
- Administration tab
 - for configuration administration 20
- administrative functions
 - configuration 20
- administrative settings 34
- alarms
 - current 19
 - types of 19
- alerts
 - transmission of 5
- array
 - firmware upgrades to 146
- array components, overview of 8, 13
- array configuration
 - resetting 157
- array hot spares 7
- array performance
 - factors in 64
 - RAID levels and 107
- array segments, storage profiles and 113
- array trays, number of 13
- arrays
 - auto-discovering 32
 - changing settings for 36
 - deleting 33
 - displaying information on 31
 - managing 31
 - methods for adding 31
 - monitoring the performance of 45
 - registering 32

- upgrading firmware 33

- asterisk, meaning of
 - in forms 22, 23
 - in searches 24

B

- banner, described 18
- boot-up failure, responding 157
- browser interface
 - navigating 18

C

- clustering software 6
- column sort order, changing 21
- configuration jobs
 - accessing 20
- console, accessing 18
- content pane, elements in 20
- controller
 - resetting 159
- controller ownership 60
- controller tray
 - described 132
 - overview of 9, 14
- controller tray LEDs and switches 148
- controllers
 - displaying performance statistics on 133
 - LEDs on 147
 - redundant 9, 14
- copy pairs
 - described 62

- recopying volume copies for 67
- removing 69
- copy priorities
 - changing 68
- copy priorities, for volume copies 64
- current alarms 19
- current user logins 19

D

- data host software 5
- data hosts
 - management software and 4
 - multipathing and 30
- data replication
 - enabling 98
- default storage configuration 50
- Default storage domain, initiators and 118
- Default storage profile, characteristics of 114
- diagnostic software 5
- disk drives 51
 - data mirroring on 8, 13
 - described 133
 - displaying information for 136
 - number of 132
 - RAID levels and 106
 - types of 132
- disk mirroring, RAID levels and 106
- disk scrubbing 36
- disk types
 - profiles and 113
- display filters, applying 21

E

- error
 - detection 36
 - reporting 36
- error detection, enabling 36
- events
 - displaying information about 44
- expansion tray 12
- expansion trays
 - described 132
 - number of 8, 13
 - overview of 15

F

- failure handling considerations, in snapshot planning 72
- fan assembly 8, 14
- fault tolerance, RAID levels and 107
- FC drives 15
- Fibre Channel (FC) ports
 - initiators and 118
- field-replaceable units (FRUs) 147
- filters, display, applying 21
- firmware upgrades 146
 - backing out of 147

G

- guest role 40

H

- hardware failure, recognition of 7
- hardware overview
 - controller tray LEDs and switches in 148
 - expansion trays in 12
- Help button 19
- help function, using 25
- High Performance Computing storage profile 114
- High_Capacity_Computing profile 114
- high-availability (HA) features 7
- host bus adapters (HBAs) 51
 - initiators and 118
 - overview of 30
- host groups
 - adding and removing members 123
 - creating 122
 - deleting 124
 - described 4
 - displaying information about 122
 - managing 122
 - mapping a volume to 58
 - mapping snapshots to 79
 - mapping to a volume 125
 - overview of 121
 - planning volume mapping for 55
 - snapshot planning and 73
 - storage domains and 139
 - unmapping 126
 - unmapping a snapshot from 80
 - unmapping a volume from 60

hosts

- access to multiple volumes by 121
- creating 128
- deleting 129
- displaying information about 127
- initiators and 118
- managing 127
- mapping a volume to 58
- mapping snapshots to 79
- mapping to a volume 129
- overview of 127
- planning volume mapping for 55
- snapshot planning and 73
- storage domains and 139
- unmapping 130
- unmapping a snapshot from 80
- unmapping a volume from 60

hot-spares, types of 133

I

initiators 51

- deleting 120
- displaying information for 118
- host bus adapters (HBAs) and 30
- mapping to hosts or host groups 55, 73
- overview of 118
- restrictions on deleting 120
- storage domains and 118, 139
- volumes and 54

interface, navigating 18

J

Java Web Console, accessing 18

jobs

- canceling 46
- displaying information for 46
- overview of 45

Jobs tab

- for configuration jobs 20

L

LED indicators

- checking 156
- described 147

LEDs

- controller tray (back) 150
- controller tray (front) 148

licenses

- adding 43
- disabling 43
- displaying information about 42
- obtaining 42

link rate LED check 156

log files 5, 7

- viewing 157

Log Out button 18

logging out 26

logical storage elements 51, 52

- configuring 20
- searching for 24

Logical Storage tab 20

logins

- current user 19

LUN masking 118

luxadm command 30

M

Mail Spool storage profile 114

management host

- isolated array and 29
- remote CLI client and 29
- system IP address and 39

management software 4

- hardware failures and 7

monitoring software 5

multipathing 30

N

navigational tabs 19

NFS Mirroring storage profile 114

NFS Striping storage profile 114

notifications

- configuring 44

number of disks, setting in storage profiles 114

O

operating system upgrades, backing out of 147

Oracle DSS storage profile 114

Oracle OLTP HA storage profile 114

Oracle OLTP storage profile 114

P

- page banner, described 18
- page display controls 21
- parity checks, RAID levels and 106
- passwords
 - changing 34
- patch upgrades 146
 - backing out of 147
- performance statistics
 - controller, displaying 133
 - volume, displaying 57
- performance, monitoring 45
- physical storage elements 50, 51
 - configuring 20
 - searching for 24
- Physical Storage tab 20
- power supplies 8, 14
- power-on procedures
 - checking the link rate LED for 156
- preferred controller 60, 158
- premium features
 - obtaining licenses for 42
- product overview
 - hardware
 - controller tray LEDs and switches 148
 - expansion tray 12

R

- RAID controllers 7, 9, 14
 - tray types and 132
- RAID levels
 - described 106
- RAID sets 4, 52, 106
- Random 1 storage profile 114
- read-ahead mode, profiles and 113
- redistributing volumes 158
- Refresh button 18
- remote CLI client
 - described 5
 - installing 29
- replicated volumes
 - described 53
- replication sets
 - described 4
- reserve volumes

- deleted snapshots and 71
- described 54, 70
- disabled snapshots and 71
- planning the capacity of 73
- threshold for 71

- Reset button 23, 24
- resetting a controller 159
- resetting an array configuration 157

S

- SATA drives 15
- Save button 23
- search function 18
 - in help 25
 - using 24
- security considerations 51
- site LAN
 - isolating the array from 29
 - managing the array from 28
- snapshot volumes, described 53
- snapshots
 - described 4
 - expanding the capacity of 80
 - mapping of host groups to 121
 - storage domains and 140
- software version, displaying 18
- source volumes
 - described 53
 - displaying 66
 - removing copy pairs and 69
 - types of 63
 - volume copies and 62
- sscs command 29
- standard volumes, described 53
- storage area networks (SANs)
 - initiators in 118
- Storage Automated Diagnostic Environment 5, 7
- storage customers, authentication of 118
- storage domains 139
 - enabling 141
 - host groups and 121
 - hosts and 127
 - licensing 141
 - methods for creating 140
 - preparing to create 140
- storage elements

- logical 52
- physical 51
- storage pools 52
 - creating 111
 - default 50
 - deleting 112
 - described 4, 110
 - displaying information for 110
 - effect of deleted volumes on 61
 - modifying 112
 - multipathing and 30
 - planning 55
 - profiles and 113
 - restrictions on deleting 110
 - storage profiles and 110
- storage profiles
 - copying 116
 - creating 115
 - deleting 117
 - described 4
 - displaying information for 115
 - modifying 116
 - optimizing attributes with 51
 - overview of 113
 - standard 114
- storage provisioning 50
- storage role 40
- storage trays 51
 - displaying information for 135
 - number of disk drives in 132
 - types of 132
- storage volumes 52
 - changing the name or description of 57
 - creating 57
 - deleting 61
 - described 4
 - displaying information for 56
 - extending the capacity of 59
 - managing 56
 - mapping to hosts or host groups 58
 - maximum number of 54
 - overview of 53
 - planning 54
 - resnapping 78
 - restrictions on deleting 61
 - storage pools and 110
 - types of 53
 - unmapping from hosts or host groups 60

- Sun Fire cabinet 8, 13
- Sun Rack cabinets 8, 13
- Sun StorageTek Enterprise Storage Manager software 6
- Sun StorageTek Expansion cabinet 8, 13
- Sun StorageTek SAM-FS software 6
- Sun StorageTek SAN Foundation Software 5
- Sun StorageTek SAN Foundation software 30
- Sun StorageTek Traffic Manager software 5, 30
- Sybase DSS storage profile 114
- Sybase OLTP HA storage profile 114
- Sybase OLTP storage profile 114
- system defaults 50
- system IP address
 - management host connection and 39
 - specifying the method for setting 38
- system name, displayed in banner 19

T

- table columns, changing the sort order of 21
- tables, filtering information in 21
- tabs, navigational 19
- target volumes
 - described 53
 - displaying 66
 - removing copy pairs and 69
 - types of 64
 - volume copies and 62
- troubleshooting 146

U

- user accounts 40
- user name, displayed in banner 19
- users
 - adding 41
 - deleting 41
 - displaying information about 40

V

- VERITAS Volume Manager with DMP 6, 30
- Version button 18
- virtual disks 4, 52, 106
 - defragmenting 108
 - displaying information for 107
 - expanding the size of 107

- planning the assignment of volumes to 55
- requirements for 72
- storage pools and 53
- volumes and 53
- volume capacity, extending 59
- volume copies
 - changing the copy priority for 68
 - creating 66
 - described 4
 - displaying information on 66
 - enabling 65
 - managing 65
 - overview of 62
 - planning 63
 - recopying 67
 - status of 63
- volume snapshots 52
 - creating 66, 77, 78
 - deleting 82
 - disabled 71
 - disabling 81
 - displaying information for 66, 68, 77
 - managing 65, 76
 - mapping to hosts and host groups 79
 - overview of 70
 - planning 72
 - resnapping 78
 - unmapping 80
 - uses for 70
 - volume capacity and 59
 - write operations and 71
- volumes
 - displaying performance statistics on 57
 - mapping of host groups to 121
 - mapping of hosts to 127
 - redistributing 158
 - storage domains and 139

W

- web browser
 - content pane elements in 20
 - navigating 18
- wildcard character, in searches 24