



Sun StorEdge™ Traffic Manager Installation and Configuration Guide

For the Solaris Operating System and the Sun
StorEdge SAN Foundation Software 4.4

Sun Microsystems, Inc.
www.sun.com

Part No. 817-3674-12
May 2005

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

Copyright 2004 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, Sun StorEdge, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

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

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

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

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

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

Copyright 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

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

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

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

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

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

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

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

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

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, 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.



Adobe PostScript

Contents

Preface	vii
1. Sun StorEdge Traffic Manager Overview and Architecture	1
Task Summary To Install and Configure the Software	2
STMS Architecture	3
Architecture Overview	3
STMS and the Solaris Operating Environment	4
Device Management Configurations	6
Summary of Features	9
The Multipath Solution	10
Multipath Configuration Management	10
Single Instance Multipath Devices	10
I/O Load Balancing	11
Failover Support	11
Dynamic Reconfiguration	11
STMS Boot	11
2. Installing the Sun StorEdge Traffic Manager Software	13
Finding and Downloading Packages and Patches	13

3. Configuring the Sun Storage Traffic Manager Software	15
Before You Begin	16
Configuring Ports Globally Through STMS	16
▼ To Enable Sun StorEdge Traffic Manager Globally	17
Configuring the Software on a Per Port (pHCI) Basis	19
Important Notes Before You Begin	20
▼ To Enable or Disable a Port	21
▼ To Configure a Single PCI HBA	23
▼ To Configure a Dual PCI HBA	25
▼ To Configure an SBus HBA	26
Configuring Third-Party Symmetric Storage Devices	27
Information You Need Before You Begin	27
▼ To Configure Third-Party Devices for Multipathing	28
▼ To Verify Third-Party Storage Devices Are Under STMS Control	29
Configuring Boot Controller Devices with the <code>stmsboot</code> command	30
Important Notes Before You Begin	32
▼ To Display Potential <code>/dev</code> Device Name Changes (from non-STMS to STMS) for All Controllers	33
▼ To Display Potential <code>/dev</code> Device Name Changes (from non-STMS to STMS) for an Individual Controller	34
▼ To Enable STMS on the Boot Controller Port	35
▼ To Disable STMS on the Boot Controller Port	38
Configuring Automatic Failback	41
▼ To Configure Automatic Failback Capability	41
▼ To Disable Automatic Failback Capability	42
Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs	43

Configuring Solstice DiskSuite or Solaris Volume Manager with STMS	45
Configuration Overview	45
▼ To Determine Pre-STMS Device Path Names	46
▼ To Unconfigure Solstice DiskSuite or Solaris Volume Manager for Devices Under STMS Control	46
▼ To Determine STMS Device Name to Pre-STMS Device Name Mapping	46
▼ To Enable STMS	46
▼ To Reconfigure the Solaris DiskSuite or Solaris Volume Manager	46
Example of Mirrored Devices	47
4. Deployment Scenarios	55
About Traffic Manager Configuration Samples	56
Configuration without Traffic Manager	57
Utility: format	57
Configuration with Traffic Manager	58
Utility: format	58
Utility: luxadm probe	59
Utility: luxadm display	59
Utility: luxadm failover	64
Example of Failover:	65
5. Troubleshooting STMS Software	69
System Crashes During STMS Boot Operations	70
System Crashes During or After an STMS Boot Enable Operation	71
System Crashes During or After an STMS Boot Disable Operation	73
STMS Not Running Properly	75
STMS Running Properly, but luxadm display and luxadm failover are Failing	75
Sun StorEdge T3, 6x20, or 3900 Arrays Do Not Show	76
System Failed During Boot with <code>scsi_vhci</code> Attachment	76

Group of Sun StorEdge A5200 Arrays Connected, But They Still Appear Under Physical Paths in Format	77
System and Error Messages	78

Preface

The intended audience for this book is system administrators who are responsible for installing the server-side software, as well as adding and configuring new workstations.

Before You Read This Book

For late breaking information, see the *Sun StorEdge SAN Foundation Software 4.4 Release Notes, Including Sun StorEdge Traffic Manager for the Solaris Operating System* (part number 817-3673-12).

How This Book Is Organized

This book contains the following sections:

Chapter 1 describes the architecture of the Sun StorEdge™ Traffic Manager software.

Chapter 2 describes information about installing the Sun StorEdge Traffic Manager software.

Chapter 3 describes how to configure the Sun StorEdge Traffic Manager software.

Chapter 4 describes deployment scenarios.

Chapter 5 describes issues you may encounter that may require troubleshooting.

There is also a Glossary containing a list of words and phrases and their definitions.

Using UNIX Commands

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

See one or more of the following for this information:

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

Typographic Conventions

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

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

Shell Prompts

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

Related Documentation

Application	Title	Part Number
Late Breaking	<i>Sun StorEdgeSAN Foundation Software Release Notes</i>	817-3673-12
SAN software	<i>Sun StorEdge SAN Foundation Software 4.4 Guide to Documentation</i>	817-3670-11
Array configuration	<i>Sun StorEdge T3+ Array Administrator's Manual</i>	816-4770
	<i>Sun StorEdge 3900 and 6900 Series 1.1 Reference and Service Manual</i>	816-5253-10
	<i>Sun StorEdge 3900 and 6900 Series 2.0 Reference and Service Manual</i>	816-5253-11
	<i>Storage Automated Diagnostic Environment User's Guide</i>	816-3142
	<i>Sun StorEdge 6020 and 6120 Arrays System Manual</i>	817-0200
	<i>Sun StorEdge 6320 System 1.0 Reference and Service Manual</i>	817-7879
	<i>Sun StorEdge 6000 Family Host Installation Software Guide</i>	817-1739

Accessing Sun Documentation

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

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

Contacting Sun Technical Support

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

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

Sun Welcomes Your Comments

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

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

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

Sun StorEdge Traffic Manager Installation and Configuration Guide: For the Solaris OS and SFS 4.4, part number 817-3674-11

Sun StorEdge Traffic Manager Overview and Architecture

This chapter describes the architecture as well as the features of Sun StorEdge™ Traffic Manager software.

Topics in this chapter include:

- “Task Summary To Install and Configure the Software” on page 2
- “STMS Architecture” on page 3
- “Summary of Features” on page 9

Task Summary To Install and Configure the Software

The *Sun StorEdge SAN Foundation Software 4.4 Guide to Documentation* lists the documentation available for this product.

TABLE 1-1 Installation and Configuration Task Summary

Task	See This Section or Manual for More Information
1. Download and install the Sun StorEdge SAN Foundation software	<i>Sun StorEdge SAN Foundation Software 4.4 Installation Guide</i> Chapter 2
2. Configure and enable the software for: <ul style="list-style-type: none">• All ports globally• Individual ports• Third-party devices	Chapter 3 “Configuring Ports Globally Through STMS” on page 17 “Configuring the Software on a Per Port (pHCI) Basis” on page 20 “Configuring Third-Party Symmetric Storage Devices” on page 28
3. Configure boot controller devices.	“Configuring Boot Controller Devices with the <code>stmsboot</code> command” on page 31
4. Configure other options.	“Configuring Automatic Failback” on page 42 “Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs” on page 44 “Configuring Solstice DiskSuite or Solaris Volume Manager with STMS” on page 46

STMS Architecture

The Sun StorEdge Traffic Manager software (STMS) is a client that manages devices that are accessed by a host through multiple paths.

STMS implements a set of interfaces for host controller interface (HCI) drivers, known as the multipath driver interface (MDI). The software provides a set of services for multipath configuration management.

This section includes the following topics:

- “Architecture Overview” on page 3
- “STMS and the Solaris Operating Environment” on page 4
- “Device Management Configurations” on page 6

Architecture Overview

The fundamental change in the STMS architecture is the reconstruction of the device tree to represent a multipath device (client) as a *single device instance* in the Solaris™ operating environment instead of *multiple instances* to represent all physical paths to the devices, as implemented in the current architecture. FIGURE 1-1 and FIGURE 1-2 show the current and the new Solaris device tree representations.

Virtual host controller interface (vHCI) drivers are pseudo nexus drivers which implement multipath services for a specific command set or bus architecture. A single instance of a vHCI driver exists for each command set which supports multipath devices.

vHCI drivers are provided with naming and transport services by one or more physical Host Controller Interface (pHCI) devices which share the common command set or bus architecture such as SCSI-3.

In the STMS architecture, an HBA is called pHCI. Also the state of a path is indicated as either ONLINE, OFFLINE, STANDBY. The paths are now classified in terms of preference.

The principal feature provided by the vHCI nexus is a path-independent bus nexus under which multipath client devices can be attached. Client devices are created as they are registered with the framework by the pHCI devices. This provides a path independent name in both `/dev` and `/devices` directories.

STMS delivers an implementation of a vHCI driver named `scsi_vhci` for SCSI-3 Fibre Channel devices. The name of the node in the Solaris device tree is:

```
/devices/scsi_vhci
```

with client devices (targets) having names of the form:

```
/devices/scsi_vhci/ssd@g200000203709C3F5:a
```

STMS and the Solaris Operating Environment

FIGURE 1-1 and FIGURE 1-2 illustrate the current and the new Solaris device trees respectively. The illustration shows a system with dual-pathed Fibre Channel Sun StorEdge A5200 storage arrays.

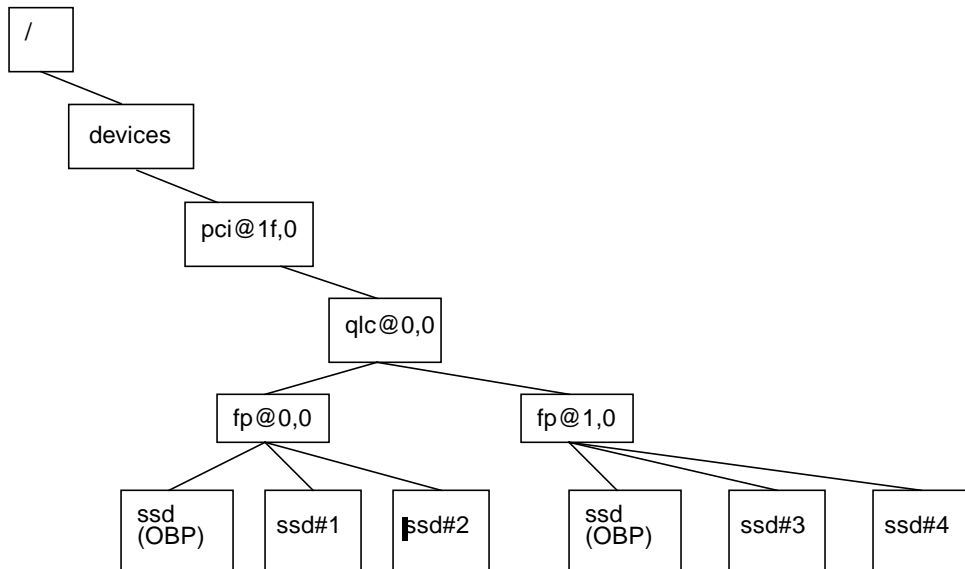


FIGURE 1-1 Current Device Tree With a Multipath Device (STMS Disabled)

In this case:

- `ssd#1` is device `ssd@w210000203709C3f5,0`
- `ssd#2` is device `ssd@w210000203709D3f5,0`
- `ssd#3` is device `ssd@w220000203709C3f5,0`
- `ssd#4` is device `ssd@w220000203709D3f5,0`

Note – Two instances of the `ssd` target device exist for a single dual-pathed Fibre Channel Device.

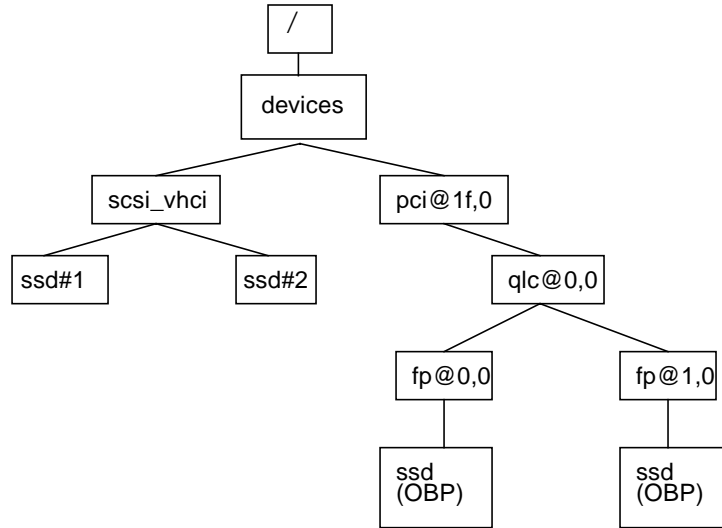


FIGURE 1-2 Subsection of Solaris Device Tree with STMS Enabled

In this case:

- `ssd#1` is device `ssd@g200000203709C3f5,0`
- `ssd#2` is device `ssd@g200000203709D3f5,0`

Note – `ssd#1` and `ssd#2` from FIGURE 1-1 have been consolidated to one `ssd` node `ssd#1` under `scsi_vhci`. Similarly `ssd#3` and `ssd#4` have been consolidated to `ssd#2` under `scsi_vhci`.

Device Management Configurations

FIGURE 1-3 shows a device management system without incorporating the Sun StorEdge Traffic Manager software. The physical device interconnects c2 through c5 connect storage devices through multiple paths to the corresponding host.

In this configuration, the format of the path from the host to a device node is "cxt_yd_z," where *x* is a unique *controller number*, *y* is the *target id* that contacts the device, and *z* is the *LUN number* associated with that device. For example, the path that refers the host to LUN 0 device node is c2t3d0.

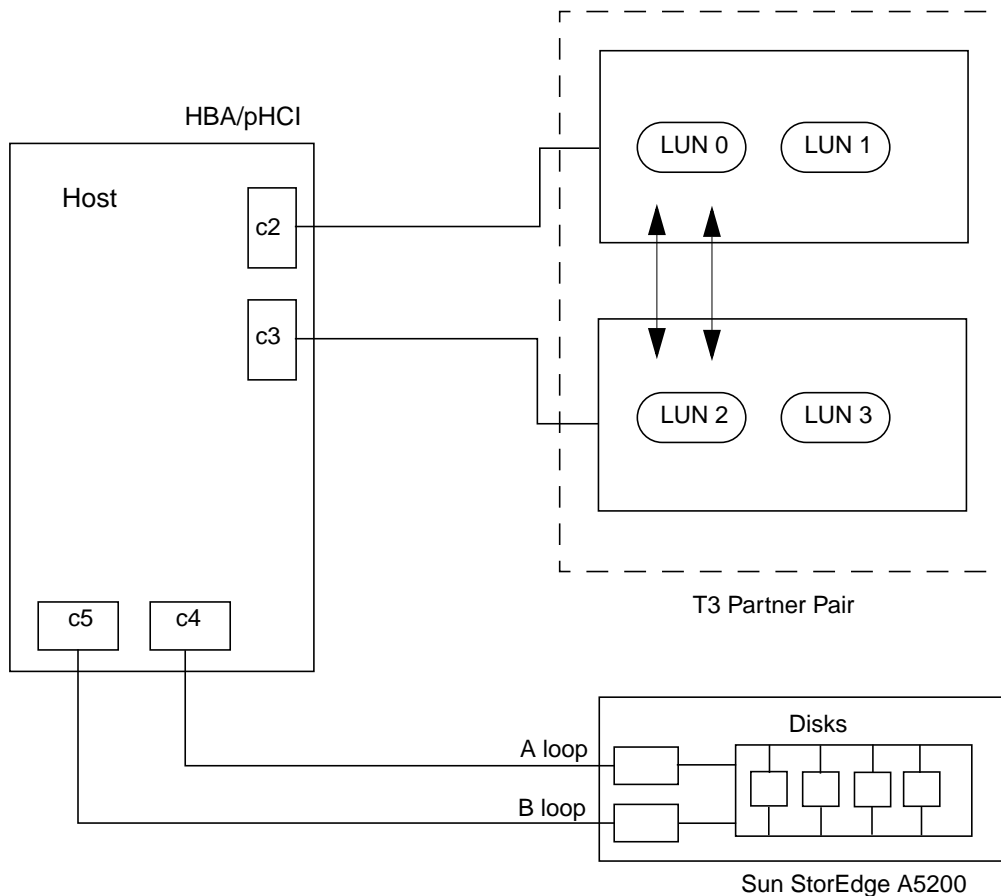


FIGURE 1-3 Device Management Configuration Without Sun StorEdge Traffic Manager Software

FIGURE 1-4 shows a device management system incorporating the Sun StorEdge Traffic Manager software

Compared to FIGURE 1-3, FIGURE 1-4 shows that only one virtual controller exists for all devices that are managed by STMS. The virtual path entry hides the complexity of multiple paths to the same device.

In this configuration, the format of the path from the host to a device node is “cxtGUIDd0,” where *x* is a *virtual controller number* and GUID is the Global Unique Identifier for the device. The GUID is a long name that uniquely identifies each Sun LUN.

A GUID is globally unique for each LUN. For example, the device names of the two LUNs of a Sun StorEdge array will have a long name as follows:

```
c6t60020F20000002E63EAF9A000022560d0
```

```
c6t60020F20000002E63EAF9BB000B8A16d0
```

You can use the `luxadm display` command for details about a device or LUN, as well as its associated multiple paths and status.




FIGURE 1-4 Device Management Configuration *With* Sun StorEdge Traffic Manager Software

Summary of Features

Note – “STMS Architecture” on page 3 describes more about multipathing, the physical and virtual host controller interfaces, and device management configurations.

The Sun StorEdge Traffic Manager software includes the following features:

- “The Multipath Solution” on page 10
- “Multipath Configuration Management” on page 10
- “I/O Load Balancing” on page 11
- “Failover Support” on page 11
- “Single Instance Multipath Devices” on page 10
- “Dynamic Reconfiguration” on page 11
- “STMS Boot” on page 11

The Multipath Solution

Modern high performance I/O bus architectures are migrating from a host-centric model where storage is private to a single host to the storage area network (SAN) model where storage is treated as a peer to host computers. In a SAN architecture, the storage is managed as a pool of resources which is shared among multiple heterogeneous hosts via a shared I/O interconnect. Combined with this model is the new requirement to provide reliability, availability, and serviceability (RAS) performance metrics, as the Solaris™ operating environment attempts to better support enterprise and high availability (HA) environments.

To increase both availability and I/O bandwidth to the storage pools, the Solaris operating environment platforms must attach to these SANs by using multiple host controller interfaces (HCIs) and I/O interconnects controller interfaces. With the addition of STMS, the Solaris operating environment is now capable of supporting the multipath device configurations presented by these new SAN architectures.

Multipath Configuration Management

The STMS architecture enables mapping of available paths to client devices to be managed by the framework undertaken by the physical host controller interface (pHCI) devices. The virtual host controller interface (vHCI) driver is also expected to supply interfaces for user-level system management applications to query and manage the pathset configurations being maintained by an instance of a vHCI device.

Single Instance Multipath Devices

The principal feature provided by the vHCI nexus is a path-independent bus nexus under which multipath client devices can be attached. Client devices are created as they are registered with the framework by the pHCI devices.

This provides both a path-independent name in `/dev` and `/devices` directories as well as eliminating the need for layered drivers which consolidates multiple device instances into a pseudo-single instance device.

I/O Load Balancing

The vHCI driver selects and routes I/O requests from its attached client devices to the pHCI device that provides the best transport services to the device. This routing decision needs to consider both the default pathset assigned to client device request as well as any routing policy such as round robin.

Failover Support

The vHCI and pHCI drivers are responsible for managing failovers, both automatic and manual. An example of an automatic failover is a cable that is unintentionally disconnected; an example of a manual failover is an administrator simulating disaster recovery by disabling a system board containing a pHCI device. Failover semantics need to be compatible with the support that is provided by the existing layered products. If an interconnect or device error is noted by a pHCI driver instance, the vHCI is notified of the loss of transport service provided by the pHCI.

Dynamic Reconfiguration

STMS supports dynamic reconfiguration.

STMS Boot

Note – The STMS Boot feature is applicable only when the boot device is connected to Sun-supported Fibre Channel HBAs and storage supported by STMS. You cannot use this feature if the boot device is a SCSI or an IDE device.

New in this release, the boot controller port can also have STMS enabled. Devices connected to the boot controller port can be multipathed client devices attached to the vHCI driver.

The Solaris 8 and 9 operating environments do not include a facility to boot a device using STMS. However, the `stmsboot` command is included with the Sun StorEdge SAN Foundation software (SFS) `SUNWmdi` package to support this feature known as STMS Boot. The `stmsboot` command provides options to enable and disable the devices connected to boot controller port. You can run this command only when STMS is enabled on the system. Devices connected to the boot controller port can be multipathed client devices attached to the virtual host control (vHCI, `scsi_vhci`) driver.

Also, applications might be using non-STMS paths for the devices connected to the boot controller port. To help ease the transition from non-STMS to STMS device names, `stmsboot` can display the STMS device name corresponding to the non-STMS paths.

See the `stmsboot(1M)` manpage for a description of the command syntax. See also “Configuring Boot Controller Devices with the `stmsboot` command” on page 31.

Installing the Sun StorEdge Traffic Manager Software

Before installing the software, ensure that your system meets the requirements identified in the *Sun StorEdge Traffic Manager Release Notes: For the Solaris OS and SFS 4.4*.

The STMS software package named `SUNWmdi` is included as part of the compressed Sun StorEdge SAN Foundation 4.4 software package files available at the <http://www.sun.com/storage/san> web site. One compressed file contains the packages for the Solaris 8 operating environment, another compressed file contains the packages for the Solaris 9 operation environment.

You can install the software using the `install_it` script or installing each package manually. For installation information, see the *Sun StorEdge SAN Foundation Software 4.4 Installation Guide* listed in the *Sun StorEdge SAN Foundation Software 4.4 Guide to Documentation*, “Related Documentation” on page x.

Finding and Downloading Packages and Patches

TABLE 2-1 lists where to find the software.

TABLE 2-1 Packages and Patches Download Web Sites

Software	Downloading Web Site
Packages	http://www.sun.com/storage/san
Patches	http://sunsolve.sun.com

Configuring the Sun Storage Traffic Manager Software

Configuring the Sun StorEdge Traffic Manager software (STMS) depends on how you intend to use your system. STMS can be configured to control all of your Sun supported Fibre Channel HBAs as listed in the *Sun StorEdge SAN Foundation Software 4.4 Release Notes, Including Sun StorEdge Traffic Manager for the Solaris Operating System*

STMS can also be configured to control a subset of your Sun supported HBAs and is capable of performing load balancing over the Sun supported HBAs that it controls. The global configuration file contains a switch to either turn on or off load balancing. This section describes what files and parameters are used to configure STMS.

You can now use STMS to configure multipathing for third-party symmetric storage devices in addition to Sun storage devices.

This chapter also provides the instructions for configuring STMS if your system is using the Solstice DiskSuite software, Solaris Volume Manager software, or third-party volume management and multipathing software.

Topics in this chapter include:

- “Configuring Ports Globally Through STMS” on page 17
- “Configuring the Software on a Per Port (pHCI) Basis” on page 20
- “Configuring Third-Party Symmetric Storage Devices” on page 28
- “Configuring Boot Controller Devices with the `stmsboot` command” on page 31
- “Configuring Automatic Failback” on page 42
- “Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs” on page 44
- “Configuring Solstice DiskSuite or Solaris Volume Manager with STMS” on page 46

Before You Begin

Reconfiguration Rebooting

You must perform a reconfiguration reboot:

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

whenever the following occurs:

- You change the "mpxio-disable" field in the `scsi_vhci.conf`, `qlc.conf`, or `fp.conf` files
- In non-fabric environments, you change the `mp_support` field for Sun StorEdge T3, 6x20, and 3900 arrays

Storage Device Types

For configuration purposes, there are two types of storage devices—symmetric and asymmetric.

Symmetric Storage Devices

Symmetric storage devices (arrays) are those devices where all the paths to logical devices are presented to the host as primary paths with STMS. Such storage devices generally allow I/O data flows to be performed on all paths to logical devices, regardless of where the I/O data flows are processed in the array. The Sun StorEdge 3510 FC array is an example of a symmetric storage device.

Any third-party symmetric storage device (array) can be configured by using the `scsi_vhci.conf` file that resides in the `/kernel/drv/` directory if it is not automatically recognized by Solaris. To configure third-party symmetric storage devices, see "To Configure Third-Party Devices for Multipathing" on page 29.

Asymmetric Storage Devices

Asymmetric storage devices (arrays) are those devices where paths to logical devices are presented to the host as primary and standby paths with STMS. Such devices generally allow I/O data flows to be performed on the primary path while the standby path can become online if the last primary path fails. The Sun StorEdge 6120 array is an example of an asymmetric storage device.

All Sun asymmetric storage devices are enabled by default. Due to the proprietary nature of failover mechanisms involved, third party asymmetric devices cannot be configured by using the `scsi_vhci.conf` file that resides in the `/kernel/drv/` directory. If a third-party asymmetric storage device is supported by STMS, please check the vendor instructions to configure it.

To configure multipathing with STMS for Sun storage devices, see “Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs” on page 44.

Configuring Ports Globally Through STMS

Unless you explicitly enabled or disabled the software on a specific port, the global settings in STMS will apply. (You can enable STMS to manage ports through global settings and override those settings with specific controls. See “Configuring the Software on a Per Port (pHCI) Basis” on page 20 for more information.)

▼ To Enable Sun StorEdge Traffic Manager Globally

1. Using any text editor, edit the `/kernel/drv/scsi_vhci.conf` file. Do not change the name and class definitions.

```
# Copyright 2001-2003 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
#pragma ident    "@(#)scsi_vhci.conf    1.9    03/06/16 SMI"
#
name="scsi_vhci" class="root";
#
# mpzio global enable/disable switch: setting mpzio-disable="no" will activate
# I/O multipathing; setting mpzio-disable="yes" disables this feature (do
# not remove this property).
#
mpzio-disable="yes";
#
# Load balancing global configuration: setting load-balance="none" will cause
# all I/O to a given device (which supports multipath I/O) to occur via one
# path. Setting load-balance="round-robin" will cause each path to the device
# to be used in turn.
#
load-balance="round-robin";
#
# Force load driver to support hotplug activity (do not remove this property).
#
ddi-forceattach=1;
#
# Automatic failback configuration
# possible values are auto-failback="enable" or auto-failback="disable"
auto-failback="disable";
#
# For enabling MPxIO support for 3rd party symmetric device need an
# entry similar to following in this file. Just replace the "SUN    SENA"
# part with the Vendor ID/Product ID for the device, exactly as reported by
# Inquiry cmd.
#
# device-type-scsi-options-list =
# "SUN    SENA", "symmetric-option";
#
# symmetric-option = 0x1000000;
```

2. To enable STMS globally, change the value of `mpxio-disable` to "no".

By default `mpxio-disable` is set to "yes" which means STMS is disabled.

3. (Optional) To enable STMS support for third-party symmetric devices, create an entry like this:

```
device-type-scsi-options-list =  
"SUN      StorEdge 3510", "symmetric-option";  
symmetric-option = 0x1000000;
```

4. If you want STMS to use all the available paths for load balancing, leave the `load-balance` field set to the default of "round-robin".

Otherwise, change the definition to "none".

5. Perform one of the following steps.

- If you want to enable STMS on the boot controller port, see "To Enable STMS on the Boot Controller Port" on page 36.
- If you do not want to enable STMS on the boot controller port at this time, perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure  
# shutdown -g0 -y -i6
```

Configuring the Software on a Per Port (pHCI) Basis

Though the Sun StorEdge Traffic Manager software might be globally enabled (as described in “To Enable Sun StorEdge Traffic Manager Globally” on page 18), individual, per-port STMS settings have priority over global settings:

- *If globally enabled*, the Sun StorEdge Traffic Manager software will not take control of a port if it is explicitly disabled as described in this section.
- Those ports not individually configured are controlled by the *global* Sun StorEdge Traffic Manager software configuration settings.

This section contains the following procedures:

- “Important Notes Before You Begin” on page 21
- “To Enable or Disable a Port” on page 22

These sections describe an example configuration procedure.

- “To Configure a Single PCI HBA” on page 24
- “To Configure a Dual PCI HBA” on page 26
- “To Configure an SBus HBA” on page 27

Important Notes Before You Begin

Before you start configuring the software by port, consider the following:

- For the Sun StorEdge Traffic Manager software to take effect, it must first be globally enabled.
- Load balancing is controlled by the global `load-balance=` variable and cannot be applied on a per-port basis.
- Volume managers that do not support the Sun StorEdge Traffic Manager software must have STMS disabled on their HBA ports.
- Devices with STMS enabled are enumerated under `/devices/scsi_vhci`. Devices with STMS disabled are enumerated under physical path names.
- With the STMS software installed, all paths to storage devices must have STMS enabled or disabled.
- Configuring STMS by port enables STMS to co-exist with other multipathing solutions like Alternate Pathing (AP), VERITAS Dynamic Multipathing (DMP), or EMC PowerPath. However:
 - Storage devices and paths should not be shared between STMS and other multipathing solutions
 - Do not mix any other multipathing software with STMS to the same storage device.
 - Partial configuration of some paths to a storage device with STMS and other paths to the same storage device with other multipathing software might cause system problems.

▼ To Enable or Disable a Port

Note – See “Important Notes Before You Begin” on page 21.

1. Log in as superuser.

Determine the HBAs that you want the Sun StorEdge Traffic Manager software to control. For example, to select the correct device, perform an `ls -l` command on `/dev/fc`. The following is an example of the command output.

```
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp0 ->
../../../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp1 ->
../../../../devices/pci@7,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp0 ->
../../../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp1 ->
../../../../devices/pci@7,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp2 ->
../../../../devices/pci@a,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          49 Apr 17 18:14 fp3 ->
../../../../devices/pci@b,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          50 Apr 17 18:14 fp4 ->
../../../../devices/pci@12,2000/SUNW,qlc@2/fp@0,0:devctl
lrwxrwxrwx 1 root    root          56 Apr 17 18:14 fp5 ->
../../../../devices/pci@13,2000/pci@2/SUNW,qlc@4/fp@0,0:devctl
lrwxrwxrwx 1 root    root          56 Apr 17 18:14 fp6 ->
../../../../devices/pci@13,2000/pci@2/SUNW,qlc@5/fp@0,0:devctl
lrwxrwxrwx 1 root    root          56 Apr 17 18:14 fp7 ->
../../../../devices/sbus@7,0/SUNW,qlc@0,30400/fp@0,0:devctl
```

Note – The `fp7` entry is a SBus HBA. The `fp5` and `fp6` include two `/pci` elements. This indicates a dual PCI HBA. The rest of the entries do not have additional PCI bridges and are single PCI HBAs.

2. Open the `/kernel/drv/fp.conf` or `/kernel/drv/qlc.conf` configuration files in a text editor such as `vi(1M)`.

Note – If you have upgraded to the Sun StorEdge SAN Foundation Software 4.4 software from an earlier release, use the `fp.conf` configuration file instead of the `qlc.conf` file from earlier releases. The `fp.conf` file has a generic syntax format that is independent of the underlying bus type (SBus, PCI, and so on) and can work with STMS-supported HBAs like `qlc`, `jfca` (JNI), and so on.

The `qlc.conf` file has a different syntax format for PCI and SBus machines.

If you have an existing `qlc.conf` file, you can convert those settings for use in the `fp.conf` file. For example, the section “To Configure a Single PCI HBA” on page 24 shows the differing syntax of these two files.

3. Explicitly enable or disable an HBA.

Add the property "mpxio-disable" to the HBA definition (for example, as shown in the sections “To Configure a Single PCI HBA” on page 24 and “To Configure a Dual PCI HBA” on page 26):

- To enable STMS on the port, set "mpxio-disable" to “no”.
- To disable STMS on the port, set "mpxio-disable" to “yes”.

4. Perform one of the following steps.

- **If you want to enable STMS on the boot controller port**, see “To Enable STMS on the Boot Controller Port” on page 36.
- **If you do not want to enable STMS on the boot controller port at this time**, perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

▼ To Configure a Single PCI HBA

Note – See “Important Notes Before You Begin” on page 21.

Perform the following steps to configure the

fp0 -> ../../devices/pci@6,2000/SUNW,qlc@2/fp@0,0:devctl
HBA entry from the `ls -l` command output shown in “To Enable or Disable a Port” on page 22, which indicates a single PCI HBA. Add the following lines to the `qlc.conf` or `fp.conf` file.

If you are using the `qlc.conf` File

1. To explicitly enable STMS on this port, enter:

```
name="qlc" parent="/pci@6,2000" unit-address="2" mpxio-disable="no";
```

2. To explicitly disable STMS on this port, enter:

```
name="qlc" parent="/pci@6,2000" unit-address="2" mpxio-disable="yes";
```

If you are using the `fp.conf` File (recommended for the Sun StorEdge SAN Foundation Software versions 4.4 and later)

1. To explicitly enable STMS on this port, enter:

```
name="fp" parent="/pci@6,2000/SUNW,qlc@2" port=0 mpxio-disable="no";
```

2. To explicitly disable STMS on this port, enter:

```
name="fp" parent="/pci@6,2000/SUNW,qlc@2" port=0 mpxio-disable="yes";
```

3. Save and exit the file.

4. Perform one of the following steps.

- If you want to enable STMS on the boot controller port, see “To Enable STMS on the Boot Controller Port” on page 36.

- **If you do not want to enable STMS on the boot controller port at this time,** perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure  
# shutdown -g0 -y -i6
```

▼ To Configure a Dual PCI HBA

Note – See “Important Notes Before You Begin” on page 21.

Perform the following steps to configure the fp6 -> ../../devices/pci@13,2000/pci@2/SUNW,qlc@5/fp@0,0:devctl HBA entry from the ls -l command output shown in “To Enable or Disable a Port” on page 22, which indicates a dual PCI HBA. Add the following lines to the qlc.conf or fp.conf file.

If you are using the qlc.conf File

1. To explicitly enable STMS on this port, type:

```
name="qlc" parent="/pci@13,2000/pci@2" unit-address="5" mpxio-disable="no";
```

2. To explicitly disable STMS on this port, type:

```
name="qlc" parent="/pci@13,2000/pci@2" unit-address="5" mpxio-disable="yes";
```

If you are using the fp.conf File (recommended for the Sun StorEdge SAN Foundation Software version 4.4 and later)

1. To explicitly enable STMS on this port, type:

```
name="fp" parent="/pci@13,2000/pci@2/SUNW,qlc@5" port=0 mpxio-disable="no";
```

2. To explicitly disable STMS on this port, type:

```
name="fp" parent="/pci@13,2000/pci@2/SUNW,qlc@5" port=0 mpxio-disable="yes";
```

3. Save and exit the file.

4. Perform one of the following steps.

- If you want to enable STMS on the boot controller port, see “To Enable STMS on the Boot Controller Port” on page 36.

- If you do not want to enable STMS on the boot controller port at this time, perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

▼ To Configure an SBus HBA

Note – See “Important Notes Before You Begin” on page 21.

Perform the following steps to configure the fp7 -> ../../devices/sbus@7,0/SUNW,qlc@0,30400/fp@0,0:devctl HBA entry from the `ls -l` command output shown in “To Enable or Disable a Port” on page 22, which indicates an SBus HBA. Add the following lines to the `fp.conf` file.

1. To explicitly enable STMS on this port, type:

```
name="fp" parent="/sbus@7,0/SUNW,qlc@0,30400" port=0 mpxio-disable="no";
```

2. To explicitly disable STMS on this port, type:

```
name="fp" parent="/sbus@7,0/SUNW,qlc@0,30400" port=0 mpxio-disable="yes";
```

3. Save and exit the file.

4. Perform one of the following steps.

- If you want to enable STMS on the boot controller port, see “To Enable STMS on the Boot Controller Port” on page 36.
- If you do not want to enable STMS on the boot controller port at this time, perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

Configuring Third-Party Symmetric Storage Devices

You can now use STMS to provide multipathing and load-balancing with more third-party symmetric storage devices than in earlier releases, enabling greater device heterogeneity in your SAN or direct-connect storage environment. Formerly, STMS worked only with the Sun StorEdge storage arrays and systems.

To use the new functionality, you must edit new parameters in the `scsi_vhci.conf` file. This section contains the following procedures:

- “To Configure Third-Party Devices for Multipathing” on page 29
- “To Verify Third-Party Storage Devices Are Under STMS Control” on page 30

Information You Need Before You Begin

- Storage devices vendor ID and product ID

You can obtain the values for the storage *vendor_id* and *product_id* variables in the `scsi_vhci.conf` file by using the `format inquiry` command on your system. See the `format(1M)` man page.

▼ To Configure Third-Party Devices for Multipathing

1. Open the `scsi_vhci.conf` file in a text editor such as `vi(1M)`.
2. Enable `mpxio` and set load balancing by changing the following setting from:

```
mpxio-disable="yes";
```

to:

```
mpxio-disable="no";
```

3. Add the `vendor_id` and `product_id` properties.

See “Information You Need Before You Begin” on page 28.

The vendor ID (*v_id*) must be eight characters long. You must specify all eight characters. Trailing characters are blanks or spaces. Tabs or the tab character are not allowed.

The product ID (*prod_id*) can be up to 16 characters long. Trailing characters are blanks or spaces. Tabs or the tab character are not allowed.

```
device-type-scsi-options-list=  
"v_id prod_id ", "symmetric-option";  
symmetric-option=0x1000000;
```

You must replace the variables with appropriate values for your system. For example:

```
device-type-scsi-options-list=  
"ven-a pid_a_upto_here", "symmetric-option",  
"ven-b pid_b_upto ", "symmetric-option",  
"ven-c pid_c ", "symmetric-option";  
symmetric-option=0x1000000;
```

4. Save and exit the `scsi_vhci.conf` file.
5. Perform one of the following steps.
 - If you want to enable STMS on the boot controller port, see “To Enable STMS on the Boot Controller Port” on page 36.

- If you do not want to enable STMS on the boot controller port at this time, perform a reconfiguration reboot now to shut down and restart your system by using the following commands.

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

▼ To Verify Third-Party Storage Devices Are Under STMS Control

1. After you configure the `scsi_vhci.conf` file, confirm that the devices nodes are created under the `/devices/scsi_vhci` file:

```
# format
```

2. If the devices are not created, confirm the `vendor_id` and `product_id` fields in the `scsi_vhci.conf` file.

Configuring Boot Controller Devices with the `stmsboot` command

Note – The STMS Boot feature is applicable only when the boot device is connected to Sun-supported Fibre Channel HBAs and storage supported by STMS. You cannot use this feature if the boot device is a SCSI or an IDE device.

The `stmsboot` command allows you to enable devices under the control of the virtual controller. You can display the `/dev` links the `stmsboot` command would create before enabling them with the `-l` or `-L` options. For additional information about the `stmsboot` command, see the `stmsboot(1M)` man page. A brief description of the syntax is presented here.

Note – Do not add or remove devices from your configuration while enabling or disabling `/dev` links with the `stmsboot` command.

Syntax

```
stmsboot [-e | -d | -L | -l controller-number]
```

Option	Description
<code>-e</code>	Enable the enumeration of devices connected to the boot controller port under <code>scsi_vhci</code> (the multipathing virtual controller). Devices on the boot controller port will be enumerated under the <code>scsi_vhci</code> virtual controller and will be controlled by STMS.
<code>-d</code>	Disable the enumeration of devices connected to the boot controller port under <code>scsi_vhci</code> . Devices on the boot controller port will be enumerated directly under the physical controller and will not be controlled by STMS.
<code>-L</code>	Display the device name changes from non-STMS device names to STMS device names that would occur if STMS were enabled.
<code>-l <i>controller_number</i></code>	Display the device name changes from non-STMS device names to STMS device names that would occur for the given controller if STMS were enabled.

This section contains the following `stmsboot` procedures:

- “Important Notes Before You Begin” on page 33
- “To Display Potential /dev Device Name Changes (from non-STMS to STMS) for All Controllers” on page 34
- “To Display Potential /dev Device Name Changes (from non-STMS to STMS) for an Individual Controller” on page 35
- “To Enable STMS on the Boot Controller Port” on page 36
- “To Disable STMS on the Boot Controller Port” on page 39

Important Notes Before You Begin



Caution – After an `stmsboot enable` operation, do not remove `/dev` links of any non-STMS device names for the boot device. This includes, but not limited to, executing commands such as `devfsadm -C`. Even though the non-STMS boot device links are stale, they are needed if you want to disable the `stmsboot` features later. Removing these links might cause your system to lose the necessary non-STMS links and the system will not boot. You would then need to manually create the links by using the `mknod` command.

- Back up data on the root device.
- You must first enable the Sun StorEdge Traffic Manager software to use the `stmsboot enable` and `disable` features.
- Plan to perform an immediate reconfiguration reboot of the host after enabling or disabling STMS on a boot controller port. Do not perform other system tasks or other activities until you perform the reconfiguration reboot. See “Before You Begin” on page 16 for the reconfiguration reboot commands.
- Once enabled, the `stmsboot` feature has a dependency on the boot path. Upon successful completion of a `stmsboot -e` command, a per-port `mpxio-disable` entry for the boot controller path is added to the `/kernel/drv/fp.conf` file automatically. (Any devices attached to the selected path will be attached as STMS clients.) *Do not change the boot path manually.* Use the `stmsboot -d` command to disable a configured boot path. This command updates both the `eeprom-boot` device and the `/kernel/drv/fp.conf` file’s port-related `mpxio-disable` entry automatically. This command then allows the system to boot by using the selected path.
- If your system fails during an enabling or disabling `stmsboot` operation, your original `/etc/path_to_inst`, `/etc/vfstab`, `/etc/system`, and `/etc/dumpadm.conf` files are saved with `.sav` file name extensions in the `/var/tmp` directory. The saved files can be useful in recovering from any unexpected system crash. Chapter 5 describes system crash and other troubleshooting information.
- If any Sun StorEdge T3, 3900, or 6x20 arrays are connected to the boot controller port, modify the settings when prompted to do so by the `stmsboot` command. See “Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs” on page 44.
- After the `stmsboot enable` or `disable` operation, device names will change. Disable any volume manager that is actively using devices connected to the boot controller port and reconfigure these devices with new STMS device names. Some applications that use non-STMS device names will need to be modified.

▼ To Display Potential /dev Device Name Changes (from non-STMS to STMS) for All Controllers

- Display the potential /dev links for all controllers, using the `stmsboot L` command.

```
# stmsboot L
Version : 1.5
non-STMS device name      STMS device name

/dev/rdisk/c1t64d0        /dev/rdisk/c9t20000020371A1D48d0
/dev/rdisk/c1t65d0        /dev/rdisk/c9t20000020371A13F3d0
/dev/rdisk/c1t66d0        /dev/rdisk/c9t20000020371A15D0d0
/dev/rdisk/c1t67d0        /dev/rdisk/c9t20000020371A1605d0
/dev/rdisk/c1t68d0        /dev/rdisk/c9t20000020371A13D7d0
/dev/rdisk/c1t69d0        /dev/rdisk/c9t20000020371A1DF7d0
/dev/rdisk/c1t70d0        /dev/rdisk/c9t20000020371A1E68d0
/dev/rdisk/c1t80d0        /dev/rdisk/c9t20000020371A1D06d0
/dev/rdisk/c1t82d0        /dev/rdisk/c9t20000020371A1C63d0
/dev/rdisk/c1t83d0        /dev/rdisk/c9t20000020371A1C61d0
/dev/rdisk/c1t84d0        /dev/rdisk/c9t20000020371A1DA7d0
/dev/rdisk/c1t85d0        /dev/rdisk/c9t20000020371A18DBd0
/dev/rdisk/c1t86d0        /dev/rdisk/c9t20000020371A1498d0

/dev/rdisk/c4t64d0        /dev/rdisk/c9t20000020371A1D48d0
/dev/rdisk/c4t65d0        /dev/rdisk/c9t20000020371A13F3d0
/dev/rdisk/c4t66d0        /dev/rdisk/c9t20000020371A15D0d0
/dev/rdisk/c4t67d0        /dev/rdisk/c9t20000020371A1605d0
/dev/rdisk/c4t68d0        /dev/rdisk/c9t20000020371A13D7d0
/dev/rdisk/c4t69d0        /dev/rdisk/c9t20000020371A1DF7d0
/dev/rdisk/c4t70d0        /dev/rdisk/c9t20000020371A1E68d0
/dev/rdisk/c4t80d0        /dev/rdisk/c9t20000020371A1D06d0
/dev/rdisk/c4t82d0        /dev/rdisk/c9t20000020371A1C63d0
```

▼ To Display Potential /dev Device Name Changes (from non-STMS to STMS) for an Individual Controller

- **Display potential /dev links for one controller, using the `stmsboot 1 controller-number` command.**

For example, to display the device name change for the boot controller port, obtain the device name (*cxt.ydz*) corresponding to the root (/) directory from the `/etc/vfstab` file, where *x* is the controller number. In this example, *x* is equal to 1.

```
# stmsboot 1 1
Version: 1.5
non-STMS device name      STMS device name

/dev/rdisk/c1t64d0        /dev/rdisk/c9t20000020371A1D48d0
/dev/rdisk/c1t65d0        /dev/rdisk/c9t20000020371A13F3d0
/dev/rdisk/c1t66d0        /dev/rdisk/c9t20000020371A15D0d0
/dev/rdisk/c1t67d0        /dev/rdisk/c9t20000020371A1605d0
/dev/rdisk/c1t68d0        /dev/rdisk/c9t20000020371A13D7d0
/dev/rdisk/c1t69d0        /dev/rdisk/c9t20000020371A1DF7d0
/dev/rdisk/c1t70d0        /dev/rdisk/c9t20000020371A1E68d0
/dev/rdisk/c1t80d0        /dev/rdisk/c9t20000020371A1D06d0
/dev/rdisk/c1t82d0        /dev/rdisk/c9t20000020371A1C63d0
/dev/rdisk/c1t83d0        /dev/rdisk/c9t20000020371A1C61d0
/dev/rdisk/c1t84d0        /dev/rdisk/c9t20000020371A1DA7d0
/dev/rdisk/c1t85d0        /dev/rdisk/c9t20000020371A18DBd0
/dev/rdisk/c1t86d0        /dev/rdisk/c9t20000020371A1498d0
```

▼ To Enable STMS on the Boot Controller Port

Note – See “Important Notes Before You Begin” on page 33.

1. **Become superuser.**
2. **If STMS is not enabled, perform the procedures in “Configuring Ports Globally Through STMS” on page 17**
3. **If any applications are running on any device connected to the boot controller port, deactivate them.**

Tip – You might need to modify applications that use the boot controller port after the system is rebooted. An example is shown in “Configuring Solstice DiskSuite or Solaris Volume Manager with STMS” on page 46. To help ease the transition to newly-named boot controller devices, save the display of the device name changes for the boot controller port as shown in “To Display Potential /dev Device Name Changes (from non-STMS to STMS) for an Individual Controller” on page 35”.

4. **Verify that no devlinks or disks commands are running.**

```
# ps elf | grep devlinks
```

```
# ps elf | grep disks
```

5. **If these commands are running, wait until they are finished.**

6. Enable the STMS software on the boot controller port.

```
# stmsboot e
Version: 1.5
WARNING!!!
    Important system files will be modified.
    The devfsadm daemon will be killed.
New STMS links for devices connected to the boot controller port
will be created.

Please make sure that there are no processes running devfsadm,
devlinks, or disks from now until the system is rebooted.
This includes, but is not limited to, initiating DR events.

Do you want to continue ? [y/n] (default: n)
```

7. Type **y** to continue.

The software then prompts you to perform a reconfiguration reboot

Note – In case of an unexpected error, the `stmsboot` operation will fail, indicating an error condition, and then exit. Some device links might have been created which are useless but the system files remain unmodified. Correct the error and you can repeat the operation. See also Chapter 5 for troubleshooting information.

If Sun StorEdge T3, 3900, or 6x20 arrays are connected to the boot controller port, the software will prompt you to change the `mp_support` setting on the storage array appropriately.

```
Boot controller is c1
```

```
Listed below are the non-STMS device names and their corresponding  
new STMS device names.
```

```
/dev/rdisk/clt4d0 /dev/rdisk/c9t60020F200000024A3CDAA4BC00051BE5d0  
/dev/rdisk/clt4d1 /dev/rdisk/c9t60020F200000024A3DCBC93E00095B12d0  
/dev/rdisk/clt4d2 /dev/rdisk/c9t60020F200000024A3DCBC95700049C40d0
```

```
WARNING!!!
```

```
There is at least one Sun StorEdge T3/T3B, or 6x20,  
that is not in "mpxio" mode.
```

```
For the system to function properly, please set  
mp_support to "mpxio" mode on the device(s) now and  
press any key to continue.
```

```
A reconfiguration reboot is mandatory for system sanity.
```

```
While you may choose to complete other tasks prior to initiating  
the reboot, doing so is not recommended.
```

```
Reboot the system now ? [y/n] (default:y)
```

Note – Do not complete other tasks before rebooting your system.

8. Type **y** to perform a reconfiguration reboot of your system.
9. If necessary, modify applications to use the new STMS device names.

▼ To Disable STMS on the Boot Controller Port

Note – See “Important Notes Before You Begin” on page 33. This `stmsboot disable` operation will have no effect on your system if you did not previously perform an `stmsboot enable` operation.

This operation changes the device names from STMS to non-STMS. If any applications are running on any device connected to the boot controller port, deactivate them. You might need to modify some applications after system is rebooted to use non-STMS device names.

1. **Become superuser.**
2. **Verify that no `devlinks` or `disks` commands are running.**

```
# ps elf | grep devlinks
```

```
# ps elf | grep disks
```

3. **If these commands are running, wait until they are finished.**
4. **Disable device configuration.**

```
# stmsboot d
Version: 1.5

WARNING!!!
    Important system files will be modified.
    The devfsadmd daemon will be killed.

Please make sure that there are no processes running devfsadm,
devlinks, or disks from now on until the system is rebooted.
This includes but, is not limited to initiating DR events.

Do you want to continue ? [y/n] (default: n)
```

5. Type y to continue.

```
Please choose a device path [1 2 ]
WARNING: Devices connected to the selected device path
will be de enumerated from scsi_vhci
1) /devices/pci@1f,4000/pci@4/SUNW,qlc@5/fp@0,0
2) /devices/pci@1f,4000/pci@2/SUNW,qlc@4/fp@0,0
Choice [default: 1]:
```

Note – In case of an unexpected error, the stmsboot operation will fail, indicating an error condition, and then exit. . Correct the error and you can repeat the operation. See also Chapter 5 for troubleshooting information.

6. Type the device path and press Return.

This example shows the results for choice 2. The software then prompts you to perform a reconfiguration reboot.

Note – If Sun StorEdge T3, 3900, or 6x20 arrays are connected to the boot controller port, the software will prompt you to change the mp_support setting on the storage array appropriately.

```
Choice [default: 1]: 2

Updated eeprom boot-device to boot through the selected path

/pci@8,600000/SUNW,qlc@2/fp@0,0/disk@w21000004cf721119,0

WARNING!!!
There is at least one Sun StorEdge T3/T3B or 6x20 in "mpxio" mode.
For the sane operation of the system, please set
mp_support to "rw" (or any appropriate mode other
than "mpxio") on the device(s) now and
press any key to continue.

A reconfiguration reboot is mandatory for system sanity.
While you may choose to complete other tasks prior to initiating
the reboot, doing so is not recommended.
Reboot the system now ? [y/n] (default:y)
```

Note – Do not complete other tasks before rebooting your system.

7. Type **y** to perform a reconfiguration reboot of your system.
8. If necessary, modify applications to use the new STMS device names.

Configuring Automatic Failback

In an earlier release of STMS, you had to manually failback a primary path with the `luxadm failover primary disk_name` command. Now you can configure the `scsi_vhci.conf` file so that primary paths fail back automatically when they become available after being in an OFFLINE state for a while. This feature is available on Sun StorEdge T3+, 3900, and 6x20 arrays.

- “To Configure Automatic Failback Capability” on page 42
- “To Disable Automatic Failback Capability” on page 43

▼ To Configure Automatic Failback Capability

1. **Open the `scsi_vhci.conf` file in your text editor.**
2. **Enable the automatic failback property.**

The default setting is “disable”;

```
auto-failback = "enable";
```

3. **Save and exit the `scsi_vhci.conf` file.**

When automatic failback is enabled, the following message is printed in the `/var/adm/messages` file during boot-up:

```
/scsi_VHCI (scsi_vhci0): Auto-failback capability enabled through  
scsi_vhci.conf file.
```

4. **Perform a reconfiguration reboot so that the changes will take effect:**

```
# touch /reconfigure  
# shutdown -g0 -y -i6
```

▼ To Disable Automatic Failback Capability

1. Open the `scsi_vhci.conf` file in your text editor.
2. Disable the automatic failback property:

```
auto-failback = "disable";
```

3. Save and exit the `scsi_vhci.conf` file.
4. Perform a reconfiguration reboot so that the changes will take effect:

```
# touch /reconfigure  
# shutdown -g0 -y -i6
```

Configuring Multipathing for Sun StorEdge T3, 6x20, and 3900 LUNs

Note – Configure and enable multipathing only when the host has STMS enabled. If host has STMS disabled on a global or per port basis, any I/O to a Sun StorEdge array in “mpxio” mode will fail on such a port until `mp_support` is set to a different mode (like `rw` or `none`).

Depending on the storage array type, you can enable or disable device multipathing support by using a command-line interface (CLI) or a storage management software user interface (UI).

TABLE 3-1 shows the various methods. The multipath or failover setting is only valid when the host has STMS enabled. After configuring multipathing (if needed), ensure that you perform a reconfiguration reboot, either from the UI or Solaris command line.

- **To perform a reconfiguration reboot from the Solaris command line, type the following:**

```
# touch /reconfigure
# shutdown -g0 -y -i6
```


TABLE 3-1 Multipathing Configuration Methods for Various Arrays

Array	Method	For More Information
Sun StorEdge T3 arrays	1. Telnet to a CLI 2. Ensure that <code>mp_support</code> is set to <code>mpxio</code> by using the <code>sys mp_support mpxio</code> command	<i>Sun StorEdge T3+ Array Administrator's Manual</i>
Sun StorEdge 3900 arrays	Default factory-delivered state of the array has the <code>mp_support</code> parameter set to <code>mpxio</code> . CLI character menu - <code>/opt/SUNWsecfg/runsecfg</code> UI - Storage Automated Diagnostic Environment software	<i>Sun StorEdge 3900 and 6900 Series 1.1 Reference and Service Manual</i> 816-5253 <i>Sun StorEdge 3900 and 6900 Series 2.0 Reference and Service Manual</i>
Sun StorEdge 6120 arrays	Default factory-delivered state of the array has the <code>mp_support</code> parameter set to <code>mpxio</code> . CLI - <code>sscs(1M)</code> command UI - Sun StorEdge Configuration Services software, installed on a Solaris host. Configure the array Failover Mode setting to Explicit	<i>Sun StorEdge 6020 and 6120 Arrays System Manual</i> <i>Sun StorEdge 6000 Family Host Installation Software Guide</i> Sun StorEdge Configuration Service online help <code>sscs(1M)</code> man page
Sun StorEdge 6320 arrays	Default factory-delivered state of the array has the <code>mp_support</code> parameter set to <code>mpxio</code> . CLI - <code>sscs(1M)</code> command UI - Sun StorEdge Configuration Services software, installed on a Solaris host. Configure the array Failover Mode setting to Explicit	<i>Sun StorEdge 6320 System 1.0 Reference and Service Manual</i> <i>Sun StorEdge 6000 Family Host Installation Software Guide</i> Sun StorEdge Configuration Service online help <code>sscs(1M)</code> man page

Configuring Solstice DiskSuite or Solaris Volume Manager with STMS

Solstice DiskSuite or Solaris Volume Manager can be reconfigured to allow the Sun StorEdge Traffic Manager software control the Sun supported HBAs it uses. The disk usage information you saved prior to enabling STMS will be used. Using the `luxadm` command you can determine the mapping from the new STMS device path name to the pre-STMS device path name. This mapping will be used to reconfigure Solstice DiskSuite or Solaris Volume Manager to recognize the correct devices under the new STMS device path names.

For more information concerning Solstice DiskSuite Solaris Volume Manager, check the Solstice DiskSuite or Solaris Volume Manager online manuals at:

<http://docs.sun.com>

This section contains the following topics:

- “Configuration Overview” on page 46
- “Example of Mirrored Devices” on page 48

Configuration Overview

The following are the steps to complete the configuration process:

Note – Backing up your data is recommended.

- “To Determine Pre-STMS Device Path Names” on page 47
- “To Unconfigure Solstice DiskSuite or Solaris Volume Manager for Devices Under STMS Control” on page 47
- “To Enable STMS” on page 47
- “To Determine STMS Device Name to Pre-STMS Device Name Mapping” on page 47
- “To Reconfigure the Solaris DiskSuite or Solaris Volume Manager” on page 47

▼ To Determine Pre-STMS Device Path Names

1. Use the `metadb` command to get a list of devices used to store the volume manager configuration.
2. Use the `metastat -p` command to get the disk-to-metadevice mapping.

▼ To Unconfigure Solstice DiskSuite or Solaris Volume Manager for Devices Under STMS Control

1. Unmount the metadevices that will be under control of STMS.
2. Clear the metadevices using the `metaclear` command.
3. Clear the metadevice database using the `metadb -df` command which takes the list of disks, output from the `metadb` command in Step 1 in To Determine Pre-STMS Device Path Names, as a list of arguments.

▼ To Determine STMS Device Name to Pre-STMS Device Name Mapping

- Use the `stmsboot -L` command output to display the potential STMS path and device name changes and the `luxadm display` command to determine which pre-STMS device names are combined under each STMS device name.

▼ To Enable STMS

- See “Configuring Ports Globally Through STMS” on page 17.

▼ To Reconfigure the Solaris DiskSuite or Solaris Volume Manager

- Create metadevices using STMS path names with the `metainit` command.



Caution – For RAID5 devices, be sure to use the “-k” option to prevent initializing the disks.

Example of Mirrored Devices

The following example shows how to migrate two-way mirrored metadevices from non-STMS to STMS. In this example:

- A Sun StorEdge T3 partner pair is connected to the host and `mp_support` on Sun StorEdge T3 is set to `rw`.
- STMS is initially disabled.
- Four LUNs of equal size exist on the partner pair. Two metadb replicas exist on lun 0 (`c2t1d0`) and lun 1 (`c2t1d1`).
- `d10` and `d11` are the sub-mirror metadevices created on the lun 2 (`c2t1d2`) and lun 3 (`c2t1d3`).
- `d14` is the mirror of `d10` and `d11`.

Note – In this section, the backslash character (\) indicates a line break. Where it is used, ensure that the command is typed on a single line.

1. Check the metadevice paths.

```
% metastat -p
d10          1 2 /dev/dsk/c2t1d2s1 /dev/dsk/c2t1d2s6
d11          1 2 /dev/dsk/c2t1d3s1 /dev/dsk/c2t1d3s6
d14          -m d10 d11
```

2. Save the pre-STMS device information.

Collect and save the output of `stmsboot -L`, `format`, `metadb`, `metastat`, and `metastat -p` commands.

```
% stmsboot -L
```

See “To Display Potential /dev Device Name Changes (from non-STMS to STMS) for All Controllers” on page 34 for an example of `stmsboot -L` output.

```
% format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
     /pci@1f,4000/scsi@3/sd@0,0
  1. c2t1d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
     ssd@w50020f2300000225,0
  2. c2t1d1 <T300 cyl 34145 alt 2 hd 20 sec 128>
     /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
     ssd@w50020f2300000225,1
  3. c2t1d2 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
     /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
     ssd@w50020f2300000225,2
  4. c2t1d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
     /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/
     ssd@w50020f2300000225,3
  5. c3t2d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,0
  6. c3t2d1 <T300 cyl 34145 alt 2 hd 20 sec 128>
     /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,1
  7. c3t2d2 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
     /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,2
  8. c3t2d3 <SUN-T300-0117 cyl 34145 alt 2 hd 32 sec 128>
     /pci@1f,2000/SUNW,qlc@1/fp@0,0/ssd@w50020f23000001f6,3
Specify disk (enter its number):
```

```
% metadb
      flags          first blk      block count
a m  pc lu0        16             1034        /dev/dsk/c2t1d0s3
a    pc lu0        1050          1034        /dev/dsk/c2t1d0s3
a    pc lu0         16             1034        /dev/dsk/c2t1d1s3
a    pc lu0        1050          1034        /dev/dsk/c2t1d1s3
```

```

% metastat -p
d14 -m d10 d11 1
d10 1 2 c2t1d2s1 c2t1d2s6 -i 32b
d11 1 2 c2t1d3s1 c2t1d3s6 -i 32b

```

```

% metastat
d14: Mirror
  Submirror 0: d10
    State: Okay
  Submirror 1: d11
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 524288 blocks
d10: Submirror of d14
  State: Okay
  Size: 524288 blocks
  Stripe 0: (interlace: 32 blocks)
    Device          Start Block  Dbase State    Hot Spare
    c2t1d2s1         0           No    Okay
    c2t1d2s6         0           No    Okay
d11: Submirror of d14
  State: Okay
  Size: 524288 blocks
  Stripe 0: (interlace: 32 blocks)
    Device          Start Block  Dbase State    Hot Spare
    c2t1d3s1         0           No    Okay
    c2t1d3s6         0           No    Okay

```

3. Unconfigure the volume manager without losing the data.

a. Clear the submirrors and mirror devices.

```
% metaclear d14
d14: Mirror is cleared
% metaclear d11
d11: Contact/Stripe is cleared
% metaclear d10
d10: Contact/Stripe is cleared
```

b. Clear the meta database replicas.

```
% metadb -d -f -c 2 /dev/dsk/c2t1d0s7 /dev/dsk/c2t1d1s7
```

Remove or comment out the `/etc/lvm/md.tab` entries.

```
#      d10          1 2 /dev/dsk/c2t1d2s1 /dev/dsk/c2t1d2s6
#      d11          1 2 /dev/dsk/c2t1d3s1 /dev/dsk/c2t1d3s6
#      d14          -m d10 d11
```

4. Enable the Sun StorEdge Traffic Manager software as described in this chapter.

See “Configuring Ports Globally Through STMS” on page 17 or “Configuring the Software on a Per Port (pHCI) Basis” on page 20.

5. Determine the Sun StorEdge Traffic Manager software device name to pre-STMS device name mapping.

The output of `stmsboot -L`, `format`, `luxadm probe`, and `luxadm display` shows the Sun StorEdge Traffic Manager software paths for each device. See “To Display Potential /dev Device Name Changes (from non-STMS to STMS) for All Controllers” on page 34 for an example of `stmsboot -L` output.

```
% format
AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
     /pci@1f,4000/scsi@3/sd@0,0
  1. c4t60020F20000002253B220F99000D348Cd0 <SUN-T300-0117
     cyl 34145 alt 2 hd 32 sec 128>
     /scsi_vhci/ssd@g60020f20000002253b220f99000d348c
  2. c4t60020F20000002253B220FC000086944d0 <SUN-T300-0117
     cyl 34145 alt 2 hd 32 sec 128>
     /scsi_vhci/ssd@g60020f20000002253b220fc000086944
  3. c4t60020F20000002253B220FD400071CD8d0 <SUN-T300-0117
     cyl 34145 alt 2 hd 32 sec 128>
     /scsi_vhci/ssd@g60020f20000002253b220fd400071cd8
  4. c4t60020F20000002253B22101D00029FC9d0 <SUN-T300-0117
     cyl 34145 alt 2 hd 32 sec 128>
     /scsi_vhci/ssd@g60020f20000002253b22101d00029fc9
Specify disk (enter its number): ^D
```

```
% luxadm probe
No Network Array enclosures found in /dev/es

Found Fibre Channel device(s):
Node WWN:50020f20000001f6 Device Type:Disk device
Logical Path:/dev/rdisk/c4t60020F20000002253B22101D00029FC9d0s2
Node WWN:50020f2000000225 Device Type:Disk device
Logical Path:/dev/rdisk/c4t60020F20000002253B220FD400071CD8d0s2
Node WWN:50020f20000001f6 Device Type:Disk device
Logical Path:/dev/rdisk/c4t60020F20000002253B220FC000086944d0s2
Node WWN:50020f2000000225 Device Type:Disk device
Logical Path:/dev/rdisk/c4t60020F20000002253B220F99000D348Cd0s2

For each entry in luxadm probe, get the device path information.
```


(The following luxadm display output shows a single LUN only.)

```
% luxadm display \  
/dev/rdisk/c4t60020F20000002253B22101D00029FC9d0s2  
DEVICE PROPERTIES for disk:  
/dev/rdisk/c4t60020F20000002253B22101D00029FC9d0s2  
  Status(Port A):      O.K.  
  Status(Port B):      O.K.  
  Vendor:              SUN  
  Product ID:          T300  
  WWN(Node):           50020f20000001f6  
  WWN(Port A):         50020f23000001f6  
  WWN(Port B):         50020f2300000225  
  Revision:            0117  
  Serial Num:          054921  
  Unformatted capacity: 68294.000 MBytes  
  Write Cache:         Enabled  
  Read Cache:          Enabled  
    Minimum prefetch:  0x0  
    Maximum prefetch:  0x0  
  Device Type:         Disk device  
  Path(s):  
    /dev/rdisk/c4t60020F20000002253B22101D00029FC9d0s2  
    /devices/scsi_vhci/ssd@g60020f20000002253b22101d00029fc9:c,raw  
  Controller           /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0  
    Device Address      50020f2300000225,3  
    Class                secondary  
    State                STANDBY  
  Controller           /devices/pci@1f,2000/SUNW,qlc@1/fp@0,0  
    Device Address      50020f23000001f6,3  
    Class                primary  
    State                ONLINE
```

6. Correlate the non-STMS device information collected in Step 1 by matching the `luxadm display` paths output. Identify the appropriate LUN.

For example, find the `c2t1d0` device under the Sun StorEdge Traffic Manager software. `c2t1d0` is LUN 0 of the Sun StorEdge T3 partner pair as seen from the format in Step 2.

```
1. c2t1d0 <T300 cyl 34145 alt 2 hd 24 sec 128>
   /pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0/ssd@w50020f2300000225,0
```

From the `luxadm display` output of each device above, check the paths, controller and device address fields.

Path(s):

```
/dev/rdisk/c4t60020F20000002253B220F99000D348Cd0s2
/devices/scsi_vhci/ssd@g60020f20000002253b220f99000d348c:c,raw
Controller /devices/pci@1f,4000/pci@4/SUNW,qlc@4/fp@0,0
Device Address 50020f2300000225,0
```

Thus:

- `c2t1d0` is now `/dev/rdisk/c4t60020F20000002253B220F99000D348Cd0`

Alternatively, the `stmsboot -L` command output also shows the device name change mapping.

7. Reconfigure the volume manager under the Sun StorEdge Traffic Manager software by recreating `/etc/lvm/md.tab` with the new device names.

a. Recreate `metadb` database replicas with new device names.

```
% metadb -a -f -c 2 \
/dev/rdisk/c4t60020F20000002253B220F99000D348Cd0s2 \
/dev/rdisk/c4t60020F20000002253B220FC000086944d0s2
```

b. Similarly, identify the metadevices with new device names by using the `metatstat -p` command.

c. Recreate submirrors and mirror.

```
% metainit d10 1 2 /dev/dsk/ \  
c4t60020F20000002253B220FD400071C8d0s1 \  
/dev/dsk/c4t60020F20000002253B220FD400071CD8d0s6  
d10: Concat/Stripe is setup
```

```
% metainit d11 1 2 /dev/dsk/ \  
c4t60020F20000002253B22101D00029F9d0s1 \  
/dev/dsk/c4t60020F20000002253B22101D00029FC9d0s6  
d11: Concat/Stripe is setup
```

```
% metainit d14 -m d10 d11  
d14: Mirror is setup
```

8. Verify that all the data is now available in the mirror device.

d14 contains the same data as before, prior to enabling the Sun StorEdge Traffic Manager software on the same mirror.

Deployment Scenarios

This chapter shows two different types of configurations:

- A system that does not use STMS
- A system that uses STMS

Topics in this chapter include:

- “About Traffic Manager Configuration Samples” on page 56
- “Configuration without Traffic Manager” on page 57
- “Configuration with Traffic Manager” on page 58

About Traffic Manager Configuration Samples

STMS device enumeration is slightly different from legacy devices in that only one device path is shown per device regardless of the number of paths.

Storage devices such as the Sun StorEdge A3x00FC have their own multipathing solution. STMS can co-exist with such storage devices. However, these devices will not be enumerated under STMS and will work the same as if STMS were not installed.

Case Study: A host has the following storage attached.

- Sun StorEdge T3 partner pair with 4 LUNs.
- Sun StorEdge A5200 with both A and B loop connected.

See FIGURE 1-3 and FIGURE 1-4 in Chapter 1 for illustrations of the two configurations described in this chapter.

Configuration without Traffic Manager

Before STMS was installed and configured, the Sun StorEdge T3 mp_support was set to "rw" format. The output is different from the configuration that included STMS.

Utility: format

The output for a configuration that does not use STMS is as follows:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
     /pci@1f,4000/scsi@3/sd@0,
  1. c2t3d0 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,0
  2. c2t3d1 <SUN-T300-0116 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,1
  3. c2t3d2 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,2
  4. c2t3d3 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@2/fp@0,0/ssd@w50020f23000042d4,3
  5. c3t4d0 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,0
  6. c3t4d1 <SUN-T300-0116 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,1
  7. c3t4d2 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,2
  8. c3t4d3 <SUN-T300-0117 cyl 34145 alt 2 hd 24 sec 128>
     /pci@1f,4000/SUNW,qlc@4/fp@0,0/ssd@w50020f2300003fad,3
  9. c4t68d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
     /pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
     ssd@w22000020371a1862,0
.
.
.
>
Specify disk (enter its number): ^D
```

Configuration with Traffic Manager

After STMS is installed and configured, Sun StorEdge T3 mp_support is set to "mpxio". (See Chapter 2.)

Utility: format

The output of `format` looks as follows:

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
   0. c0t0d0 <SUN18G cyl 7506 alt 2 hd 19 sec 248>
      /pci@1f,4000/scsi@3/sd@0,0
   1. c6t60020F20000042D43ADCBC4E000C41E2d0 <SUN-T300-0117
      cyl 34145 alt 2 hd 24 sec 128>
      /scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2
   2. c6t60020F20000042D43B0E926A000AA3FCd0 <SUN-T300-0116
      cyl 34145 alt 2 hd 24 sec 128>
      /scsi_vhci/ssd@g60020f20000042d43b0e926a000aa3fc
   3. c6t60020F20000042D43B2753510008C9DFd0 <SUN-T300-0117
      cyl 34145 alt 2 h d 24 sec 128>
      /scsi_vhci/ssd@g60020f20000042d43b2753510008c9df
   4. c6t60020F20000042D43B275377000877DDd0 <SUN-T300-0117
      cyl 34145 alt 2 h d 24 sec 128>
      /scsi_vhci/ssd@g60020f20000042d43b275377000877dd
   5. c6t20000020371A1862d0 <SUN36G cyl 24620 alt 2 hd 27 sec
107>
      /scsi_vhci/ssd@g20000020371a1862
Specify disk (enter its number): ^D
```

Consider the following notes:

1. Devices enumerated under STMS will have a `/scsi_vhci/ssd` entry.
2. The first four `scsi_vhci` entries are the four T3 LUNs which now have a long name which is Global Unique Identifier (GUID). Also, only one entry per LUN is visible instead of two paths.
3. The next `scsi_vhci` entry is a disk in the Sun StorEdge A5200 array.
4. There is only one controller number assigned for all the devices encapsulated under STMS.

5. Use `luxadm display` to identify the mapping between device entries without STMS and the device entries with STMS.

Utility: luxadm probe

The `luxadm probe` now shows the WWN for Sun StorEdge A5200 and GUID for T3.

```
# luxadm probe
Found Enclosure:
SENA Name:f   Node WWN:50800200000777d0   Logical Path:/dev/es/ses0
Found Fibre Channel device(s):
  Node WWN:50020f20000042d4   Device Type:Disk device
    Logical Path:/dev/rdisk/c6t60020F20000042D43B275377000877DDd0s2
  Node WWN:50020f20000042d4   Device Type:Disk device
    Logical Path:/dev/rdisk/c6t60020F20000042D43B2753510008C9DFd0s2
  Node WWN:50020f20000042d4   Device Type:Disk device
    Logical Path:/dev/rdisk/c6t60020F20000042D43B0E926A000AA3FCd0s2
  Node WWN:50020f20000042d4   Device Type:Disk device
    Logical Path:/dev/rdisk/c6t60020F20000042D43ADCBC4E000C41E2d0s2
```

Utility: luxadm display

The `luxadm display` has been enhanced for STMS. For each entry in the format or `luxadm probe`, `luxadm display` output indicates the following:

1. The number of paths to the storage device.
2. The mapping of the paths prior to STMS and after STMS (under Path(s) controller and device address).
3. The state and type of each path.
 - a. State of each path.
 - i. ONLINE indicates the active path(s) through which IO is going to the device. For more than one ONLINE path, STMS will use load balancing (like round robin scheme) or single IO to the device, depending on the setting of `load-balance` variable in `/kernel/drv/scsi_vhci.conf` file.
 - ii. STANDBY indicates the path is available if an ONLINE path fails or is switched to another state. There can be many STANDBY paths. If a STANDBY path is chosen to be an active path for routing IO, its status will be changed to ONLINE.

iii. OFFLINE indicates the path(s) was previously existing but is not available now.

b. Class type for each path.

i. PRIMARY: This path is the preferred path for routing IO.

ii. SECONDARY: This path is the next priority path after PRIMARY.

If a Sun StorEdge T3 partner pair configuration is used, two paths for each lun exist. One path is ONLINE and the other is STANDBY. In this configuration, IO to the lun is active on its ONLINE path. If this path fails, the STANDBY path becomes the ONLINE path. However, if the first path later becomes available, it will go into STANDBY mode instead of ONLINE, thus saving an expensive failover operation, if auto-failback was disabled throughout the configuration. Use `luxadm failover` command to bring the restored path to ONLINE (see the `luxadm man` pages).

There can be more than one primary and secondary path to a Sun StorEdge T3, 39x0 or 6x20 device in a SAN environment."

As can be seen from the following two luxadm display output of some T3 LUNs, PRIMARY can be STANDBY and SECONDARY can be ONLINE for one LUN on the same physical paths, whereas SECONDARY is STANDBY and PRIMARY is ONLINE for another LUN.

```
# luxadm display
/dev/rdisk/c6t60020F20000042D43B275377000877DDd0s2
DEVICE PROPERTIES for disk:
/dev/rdisk/c6t60020F20000042D43B275377000877DDd0s2
  Status(Port A):      O.K.
  Status(Port B):      O.K.
  Vendor:              SUN
  Product ID:          T300
  WWN(Node):           50020f20000042d4
  WWN(Port A):         50020f2300003fad
  WWN(Port B):         50020f23000042d4
  Revision:            0117
  Serial Num:          Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:         Enabled
  Read Cache:          Enabled
    Minimum prefetch:  0x0
    Maximum prefetch:  0x0
  Device Type:         Disk device
  Path(s):
    /dev/rdisk/c6t60020F20000042D43B275377000877DDd0s2
    /devices/scsi_vhci/ssd@g60020f20000042d43b275377000877dd:c,raw
      Controller        /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
        Device Address  50020f23000042d4,3
          Class          secondary
            State        ONLINE
      Controller        /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
        Device Address  50020f2300003fad,3
          Class          primary
            State        STANDBY
```

```

# luxadm display
/dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
  Status(Port A):      O.K.
  Status(Port B):      O.K.
  Vendor:              SUN
  Product ID:          T300
  WWN(Node):           50020f20000042d4
  WWN(Port A):         50020f23000042d4
  WWN(Port B):         50020f2300003fad
  Revision:            0117
  Serial Num:          Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:         Enabled
  Read Cache:          Enabled
    Minimum prefetch:  0x0
    Maximum prefetch:  0x0
  Device Type:         Disk device
  Path(s):
    /dev/rdsk/c6t60020F20000042D43B0E926A000AA3FCd0s2
    /devices/scsi_vhci/ssd@g60020f20000042d43b0e926a000aa3fc:c,raw
  Controller           /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
    Device Address      50020f23000042d4,1
    Class               primary
    State               ONLINE
  Controller           /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
    Device Address      50020f2300003fad,1
    Class               secondary
    State               STANDBY

```

In the case of Sun StorEdge A5200 in the current configuration, there are two paths. Both are ONLINE and IO is load balanced on both the paths. If a path fails, the second path continues the IO. If the failed path comes back, it will be in ONLINE state and starts participating in IO transfer if load-balancing is enabled.

```
# luxadm display /dev/rdisk/c6t20000020371A1862d0s2
DEVICE PROPERTIES for disk: /dev/rdisk/c6t20000020371A1862d0s2
Status(Port A):      O.K.
Status(Port B):      O.K.
Vendor:              SEAGATE
Product ID:          ST136403FSUN36G
WWN(Node):           20000020371a1862
WWN(Port A):         21000020371a1862
WWN(Port B):         22000020371a1862
Revision:            114A
Serial Num:          LT0187150000
Unformatted capacity: 34732.891 MBytes
Read Cache:          Enabled
  Minimum prefetch:  0x0
  Maximum prefetch:  0xffff
Location:            In the enclosure named: f
Device Type:         Disk device
Path(s):
/dev/rdisk/c6t20000020371A1862d0s2
/devices/scsi_vhci/ssd@g20000020371a1862:c,raw
Controller           /devices/pci@1f,2000/pci@1/SUNW,qlc@4/fp@0,0
  Device Address      22000020371a1862,0
  Class               primary
                    State          ONLINE
Controller           /devices/pci@1f,2000/pci@1/SUNW,qlc@5/fp@0,0
  Device Address      21000020371a1862,0
  Class               primary
  State               ONLINE
```

Utility: luxadm failover

The `luxadm failover` command is used to failover LUNs from primary to secondary paths and vice versa.

In case of Sun StorEdge T3 partner pair, 6x20 HA, or 39x0 configuration, when a failover happens (for example, a cable is removed from one storage controller), the LUNs owned by that controller are failed over to the alternate controller.

After a failure of one path is corrected, the LUN will not fail back to the original configuration automatically if auto-failback is set to "disable" in `/kernel/drv/scsi_vhci.conf`. A `luxadm failover` subcommand must be issued to perform failover to the original configuration.

Example of Failover:

1. The original state of a Sun StorEdge T3 LUN is obtained by `luxadm display` command.

```
# luxadm display /dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
DEVICE PROPERTIES for disk:
/dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
  Status(Port A):      O.K.
  Status(Port B):      O.K.
  Vendor:              SUN
  Product ID:          T300
  WWN(Node):           50020f20000042d4
  WWN(Port A):         50020f23000042d4
  WWN(Port B):         50020f2300003fad
  Revision:            0117
  Serial Num:          Unsupported
  Unformatted capacity: 51220.500 MBytes
  Write Cache:         Enabled
  Read Cache:          Enabled
    Minimum prefetch:  0x0
    Maximum prefetch:  0x0
  Device Type:         Disk device
  Path(s):
    /dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
    /devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
      Controller        /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
        Device Address  50020f23000042d4,0
        Class           primary
        State           ONLINE
      Controller        /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
        Device Address  50020f2300003fad,0
        Class           secondary
        State           STANDBY
```

Note – The primary path is ONLINE and the secondary path is STANDBY.

2. Cable is pulled from the T3 controller 50020f23000042d4. Failover is triggered and the primary path is OFFLINE and secondary path is ONLINE. LUN status now is degraded.

```
# luxadm display /dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
DEVICE PROPERTIES for disk:
/dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
..
Path(s):
/dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
/devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
Controller          /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
  Device Address    50020f23000042d4,0
  Class             primary
  State             OFFLINE
Controller          /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
  Device Address    50020f2300003fad,0
  Class             secondary
  State             ONLINE
```

3. Cable is reinserted to T3 controller 50020f23000042d4. The device state become optimal but, the failover is not triggered. The Primary path comes up as STANDBY and still the secondary path is ONLINE.

```
# luxadm display /dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
DEVICE PROPERTIES for disk:
/dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
..
Path(s):
/dev/rdisk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
/devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
Controller          /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
  Device Address    50020f23000042d4,0
  Class             primary
  State             STANDBY
Controller          /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
  Device Address    50020f2300003fad,0
  Class             secondary
  State             ONLINE
```


4. Type the `luxadm failover primary` command to failover to the primary path.

```
# luxadm failover primary
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
#
# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
..
Path(s):
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
/devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
Controller      /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
  Device Address  50020f23000042d4,0
  Class           primary
  State           ONLINE
Controller      /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
  Device Address  50020f2300003fad,0
  Class           secondary
  State           STANDBY
```

This now triggers the failover. The primary path becomes ONLINE and the secondary path becomes STANDBY which is equivalent to the original state in Step 1.

9. To verify the failover operation, display the properties using the `luxadm display` command.

```
# luxadm display /dev/rdsk/c5t60020F20000042D43ADCBC4E000C41Ed0s2
DEVICE PROPERTIES for disk:
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
..
Path(s):
/dev/rdsk/c5t60020F20000042D43ADCBC4E000C41E2d0s2
/devices/scsi_vhci/ssd@g60020f20000042d43adcbc4e000c41e2:c,raw
Controller      /devices/pci@1f,4000/SUNW,qlc@2/fp@0,0
  Device Address  50020f23000042d4,0
  Class           primary
  State           ONLINE
Controller      /devices/pci@1f,4000/SUNW,qlc@4/fp@0,0
  Device Address  50020f2300003fad,0
  Class           secondary
  State           STANDBY
```

For more details on the `luxadm failover` command, please refer to the man pages.

Troubleshooting STMS Software

This chapter provides solutions to potential problems that may occur while running Sun StorEdge Traffic Manager software.

Topics in this chapter include:

- “System Crashes During STMS Boot Operations” on page 70
- “System Crashes During or After an STMS Boot Enable Operation” on page 71
- “System Crashes During or After an STMS Boot Disable Operation” on page 73
- “STMS Not Running Properly” on page 75
- “STMS Running Properly, but luxadm display and luxadm failover are Failing” on page 75
- “Sun StorEdge T3, 6x20, or 3900 Arrays Do Not Show” on page 76
- “System Failed During Boot with `scsi_vhci` Attachment” on page 76
- “Group of Sun StorEdge A5200 Arrays Connected, But They Still Appear Under Physical Paths in Format” on page 77
- “System and Error Messages” on page 78

System Crashes During STMS Boot Operations

If your system fails during and STMS boot enable (`stmsboot -e`) or disable (`stmsboot -d`) operation, your original `/etc/path_to_inst`, `/etc/vfstab`, `/etc/system`, and `/etc/dumpadm.conf` files are saved with `.sav` file name extensions in the `/var/tmp` directory.

Other important debug information is also saved in this directory.

System Crashes During or After an STMS Boot Enable Operation

If your system crashes or panics during or after an STMS boot enable (`stmsboot -e`) operation with messages similar to the following:

```
Cannot assemble drivers for root
/scsi_vhci/ssd@g20000004cf721119:a
Cannot mount root on /scsi_vhci/ssd@g20000004cf721119:a fstype ufs

or

cannot stat /dev/rdisk/...
```

Then check the following:

1. If the boot device is a Sun StorEdge T3 Raid Array LUN, did you change the `mp_support` setting to “`mpxio`” mode? If the change was not performed before system reboot, make this change and shut down and restart the system.
2. Are you specifying a different boot path manually? `stmsboot` has a dependency on the path through which the system boots. See “Configuring Boot Controller Devices with the `stmsboot` command” on page 31 and “Important Notes Before You Begin” on page 33.
3. Did you have a per-port `mpxio-disable` entry corresponding to the boot controller port in the `qlc.conf` file? If so, remove the corresponding `qlc.conf` entry.
4. Boot the system using another disk or over the network and mount the boot device on `/mnt`.
 - Verify that `/mnt/etc/vfstab` is the same as `/var/tmp/vfstab`. Are all the boot controller device names shown as STMS device names, compared with `/var/tmp/vfstab.sav`?
 - Verify that `/mnt/etc/path_to_inst` is the same as `/var/tmp/path_to_inst`. This file should contain a unique `ssd` instance for each `scsi_vhci/ssd@gGUID` device.
 - The `/mnt/etc/system` file should have a `rootdev` entry to the boot device of `scsi_vhci/ssd@gGUID`. For example:

```
rootdev:/scsi_vhci/ssd@g20000004cf721119:a
```

- The `/mnt/kernel/drv/fp.conf` file should have the boot controller port entry set to `mpxio-disable="no"`. See "Configuring the Software on a Per Port (pHCI) Basis" on page 20.

5. To restore the system to the "stmsboot disabled" state, perform the following commands:

```
# cp /mnt/var/tmp/vfstab.sav /mnt/etc/vfstab

# cp /mnt/var/tmp/system.sav /mnt/etc/system

# cp /mnt/var/tmp/dumpadm /mnt/etc/dumpadm.conf
```

- Note that `path_to_inst.sav` is not restored because system remembers the device names. Remove the per-port `mpxio-disable` entry corresponding to the boot controller port in the `/mnt/kernel/drv/fp.conf` file and reboot your system:

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

If none of the above tips help, contact Sun Services.

System Crashes During or After an STMS Boot Disable Operation

If your system crashes or panics during or after an STMS boot disable (`stmsboot -d`) operation, check the following:

If after STMS boot disable (`stmsboot -d`) operation, your system panics similar to the following manner:

with messages similar to the following:

```
Cannot assemble drivers for root /dev/rdisk/....  
Cannot mount root on fstype /dev/rdisk/.... ufs  
or  
  
cannot stat /dev/rdisk/...
```

Then check the following:

1. If the boot device is a Sun StorEdge T3 Raid Array LUN, did you change the `mp_support` setting to “`rw`” or “`none`”? If the change was not performed before system reboot, make this change and shut down and restart the system.
2. `stmsboot` has a dependency on the path through which the system boots. See “Configuring Boot Controller Devices with the `stmsboot` command” on page 31 and “Important Notes Before You Begin” on page 33.
3. Did you have a per-port `mpxio-disable` entry corresponding to the boot controller port in the `qlc.conf` file? If so, remove the corresponding `qlc.conf` entry.
4. Boot the system using another disk or over the network and mount the boot device on `/mnt`.
 - The `/mnt/etc/vfstab` file should indicate the non-STMS device names instead of STMS device names.
 - The `/mnt/etc/system` file should not have a `rootdev` entry.
5. The `/mnt/kernel/drv/fp.conf` file should not have the boot controller port entry.

6. To restore the system to the “stmsboot enabled” state, perform the following commands:

```
# cp /mnt/var/tmp/vfstab.sav /mnt/etc/vfstab  
  
# cp /mnt/var/tmp/system.sav /mnt/etc/system  
  
# cp /mnt/var/tmp/dumpadm /mnt/etc/dumpadm.conf
```

- Note that `path_to_inst.sav` is not restored because system remembers the device names. Add the per-port `mpxio-disable="no"` entry corresponding to the boot controller port in the `/mnt/kernel/drv/fp.conf` file and reboot your system:

```
# touch /reconfigure  
# shutdown -g0 -y -i6
```

See “Configuring the Software on a Per Port (pHCI) Basis” on page 20.

If none of the above tips help, contact Sun Services.

STMS Not Running Properly

The first item to check is if STMS has been installed correctly.

Verify STMS drivers (that is, `scsi_vhci` and `mpxio`), are loaded with the help of the `modinfo` command. (See the `modinfo` man pages.)

```
# modinfo | grep mpxio
23 102193e5 84cb - 1 mpxio (MDI Library)
# modinfo | grep scsi_vhci
121 781ea000 6a20 225 1 scsi_vhci (Sun Multiplexed SCSI vHCI)
```

If `modinfo` output does not show the drivers are loaded, verify that the following binaries have been loaded in the proper directories.

- `/kernel/drv/scsi_vhci`, `/kernel/drv/sparcv9/scsi_vhci` (64 bit)
- `/kernel/misc/mpxio`, `/kernel/misc/sparcv9/mpxio` (64-bit)
- `/kernel/drv/scsi_vhci.conf`

If these binaries are not present in the specified directories, then STMS did not install properly. Please repeat the installation process. Make sure you are logged in as superuser when installing the software.

STMS Running Properly, but `luxadm display` and `luxadm failover` are Failing

Ensure that all your qlc pHCIs are seen using the `luxadm qlgc` command. If you don't see one or more of the pHCIs, then their Fcode is out of date. An alternate way to confirm this is by examining the device path. Notice the “`scsi@4`” in the example below. Make sure that none of the device paths looks like the following:

```
/devices/pci@9d,600000/pci@1/scsi@4/fp@0,0:devctl
```

To correct this problem, download the latest Fcode software.

Sun StorEdge T3, 6x20, or 3900 Arrays Do Not Show

A Sun StorEdge T3, 6x20, or 3900 array can be correctly configured for STMS by using the latest firmware revision. Be sure to set the master settings as follows:

```
t31> sys mp_support mpzio
```

See the *Sun StorEdge Traffic Manager Release Notes: For the Solaris OS and SFS 4.4* for more details about the system requirements.

While checking the array setup, confirm that the LPC version is also current. Once the device is configured, it is advisable to perform a reconfigure boot or equivalent.

System Failed During Boot with scsi_vhci Attachment

It is most likely due to an incomplete installation. The installation failed to provide an entry for the `scsi_vhci` in the `name_to_major` database or duplicate entries for the same number with different driver name references. If you do not see any installation log errors, call support.

Group of Sun StorEdge A5200 Arrays Connected, But They Still Appear Under Physical Paths in Format

Consider the following:

- Check whether the Sun StorEdge A5200 class devices are not connected to a STMS supported HBA mentioned in Chapter “Requirements and Environments.”.
- Check if STMS is disabled under `/kernel/drv/scsi_vhci.conf`.
- Check if the system is booting from the Sun StorEdge A5200 disks (booting from StorEdge T3 is not supported). When you boot from a STMS device, all devices under the pHCI with the boot device get enumerated under `scsi_vhci`.
- Make sure the pHCI has `mpxio-disable="no"` set under STMS in the `/kernel/drv/qlc.conf` or `/kernel/drv/fp.conf` files.

System and Error Messages

The following messages might appear in the course of operation.

- When the automatic failback feature is enabled through the configuration file, you should receive the following message:

```
Auto-failback capability enabled through scsi_vhci.conf file.
```

- If automatic failback succeeds, the following message is logged:

```
Auto failback operation succeeded for devices.
```

- If automatic failback fails, the following message is logged:

```
Auto failback operation failed for device.
```

- Externally initiated failover of a Sun StorEdge T3 or T3+ array has been observed.

```
Waiting for externally initiated failover to complete
```

```
path target address %s is now ONLINE because of an externally  
initiated failover
```

```
path target address %s is now STANDBY because of an externally  
initiated failover %s (GUID %s)
```

- Failover has been prevented by the cause indicated

```
Failover operation failed for device: due to an internal error
```

```
Failover operation failed for device: Invalid path-class
```

```
Failover operation failed for device: no alternate paths found
```

```
Failover operation completed successfully for device: failed over  
from _ to _
```

```
T3 failover failed, couldn't transport packet
```

```
T3 failover failed: timed out
```

```
T3 failover failed: timed out waiting for path to become active
```

```
T3 failover failed; sense key:%x, ASC: %x, ASCQ:%x
```

- After an `stmsboot enable (stmsboot -e)` operation, the following message corresponding to the boot controller port (here, `fp2`) *should not* appear once the system is booted. (That is, booted through the `stmsboot`-enabled device through the `fp2` path):

```
mpxio: [ID 284422 kern.info] /pci@8,600000/SUNW,qlc@1/fp@0,0 (fp2)  
multipath capabilities disabled: controller for root filesystem  
cannot be multipathed.
```


Glossary

A

Array A Sun StorEdge array that contain disk drives with Fibre Channel connectivity to the data host and may contain an internal RAID controller

C

Client Device A multipath device attached beneath an instance of a vHCI driver.

F

Failover Routing I/O requests through alternate paths when a failure is detected on the primary or active HCI.

G

Global Unique Identifier (GUID) While the definition is command set specific, a GUID generally is an identifier that has been defined unique, as in no 2 devices will have the same ID.

H

- HCI** The host controller interface (HCI) consists of the software and hardware required to connect the storage array to the server.
- Hot Standby** A pHCI which is configured to not transport data traffic unless all primary interfaces have failed. A pHCI configured in this way may carry a small amount of traffic for latent fault detection to ensure that the standby is operational.

L

- Load Balancing** The ability to route I/O requests through different paths for better utilization of host controller resources.
- Logical Unit Number (LUN)** One or more drives that are grouped into a unit; they are also called a “volume” (in terms of SCSI representation).

M

- MDI** The multipath driver interface (MDI) is a new set of interfaces for HCI drivers.

P

- Packet Transport** The SCSI vHCI driver interfaces with the Traffic Manager architecture to select a suitable path for routing a request.
- Partner Group** A pair of interconnected arrays.
- Path** Another term for a pHCI transport service which is being managed by a vHCI multipathing nexus driver. The vHCI driver selects a path from a list of available paths.

Path Set Set of paths which share a common set of attributes. Path sets are created automatically as part of the device discovery and enumeration process. Path sets may also be instantiated and managed by applications.

**Physical Host
Controller Interface
(pHCI)**

(a.k.a. HCI) An instance of bus nexus device providing command set specific transport and naming services for devices attached to it via an I/O interconnect.

R

Repair Detection Detecting that a path has begun to function correctly after having previously failed.

S

SCSA Sun compliant SCSI adapter.

Switchover Routing I/O through other valid paths through an administrative utility. Similar to the Alternate Pathing disk driver developed by DHPG.

V

vHCI A “virtual” pseudo nexus driver under which multipath client devices are attached. The vHCI driver is responsible for routing transport requests from the client devices to the appropriate pHCI devices which are providing transport services to the device.

Volume One or more drives that are grouped together into a unit; they are also called “LUN” (in terms of SCSI representation).

W

**World Wide Name
(WWN)**

Unique number assigned to each device on a Fibre Channel loop.

Index

C

- configuration 19
 - examples 55, 57, 58
 - examples, luxadm command 59, 64, 65
- configuring 15
 - boot controller devices 30
 - HBAs 25, 26
 - third-party devices 27

D

- deployment 55
- disabling ports 21
- download web sites 13
- downloading software 13

E

- enabling ports 21
- enabling STSM 17
- error messages 78

P

- package and patch web sites 13
- packages 13
- patches 13

S

- STMS boot 11
- stmsboot command 30, 33, 34
 - enabling 35
- Sun StorEdge Traffic Manager
 - architecture 3
 - configuring 15, 16, 19
 - downloading the software 13
 - enable STMS boot 35
 - enabling 17
 - installation 13
 - overview 3
 - package and patch web site 13
 - STMS boot 11, 30
 - stmsboot command 30, 33, 34
- sunsolve 13

T

- troubleshooting 69

