



Solaris Express SAN Configuration and Multipathing Guide



Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 820-3070-13
August 2009

Copyright 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more U.S. patents or pending patent applications in the U.S. and in other countries.

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

This distribution may include materials developed by third parties.

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 other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, the Solaris logo, the Java Coffee Cup logo, docs.sun.com, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and 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.

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

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 2009 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuelle relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plusieurs brevets américains ou des applications de brevet en attente aux Etats-Unis et dans d'autres pays.

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

Certains composants de ce produit peuvent être dérivées du logiciel Berkeley BSD, licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays; elle est licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, le logo Solaris, le logo Java Coffee Cup, docs.sun.com, Java et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., ou ses filiales, 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.

Les produits qui font l'objet de cette publication et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes chimiques ou biologiques ou pour le nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

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

Contents

Preface	7
1 Product Overview	11
About Sun StorageTek Traffic Manager	11
What's New?	12
Software Features	12
Fibre Channel Software Features	12
SAS Software Features	13
Multipathing Software Features	14
Supported Standards	15
2 Fibre Channel Configuration Overview	17
Configuration Task Overview	17
Configuration Considerations	18
3 Administering Multipathing Devices Through <code>mpathadm</code> Commands	21
Administering Multipathing Devices	21
▼ To Get a List of Multipathing Support That Can Be Administered by the <code>mpathadm</code> Command	21
▼ To Get the Multipathing Support Properties	21
▼ To Get a List of Initiator Ports Discovered on the System	23
▼ To Get the Properties of a Particular Initiator Port	23
▼ To Get a List of Logical Units Discovered on the System	23
▼ To Get the Properties and Associated Path Information of a Particular Logical Unit by Name	24
▼ To List All Logical Units Associated With a Particular Target Port	25
▼ To List a Logical Unit With a Particular Name Derived From the Hardware Used by the System	27

▼ To Configure Auto-Failback for a Particular Multipathing Support	28
▼ To Fail Over a Logical Unit	30
▼ To Enable a Path	32
▼ To Disable a Path	34
4 Configuring Multipathing Software	37
Configuring Multipathing	37
Multipathing Considerations	38
Enabling and Disabling Multipathing	38
▼ To Enable Multipathing	39
▼ To Disable Multipathing	40
Enabling or Disabling Multipathing on a Per-Port Basis	41
Port Configuration Considerations	41
Configuring Third-Party Storage Devices	43
Third-Party Device Configuration Considerations	43
Configuring Third-Party Storage Devices: New Devices	44
Configuring Third-Party Storage Devices: Disabling Devices	45
Displaying Device Name Changes	46
Configuring Automatic Failback	47
▼ To Configure Automatic Failback	47
5 Configuring SAN Devices	49
SAN Device Considerations	49
Adding SAN Devices	50
▼ To Add a SAN Device	50
Configuring Fabric Boot Devices on SPARC	51
Boot Considerations	51
6 Configuring Virtual Fibre Channel Ports	53
What is NPIV?	53
Limitations of NPIV	53
Creating NPIV Ports	54
▼ To Create an NPIV Port	54
▼ To Delete an NPIV Port	55

▼ To Report Status on NPIV Ports	55
7 Configuring FCoE Ports	57
What is FCoE?	57
Limitations of FCoE	57
Configuring FCoE Ports	58
▼ To Create an FCoE Port	58
▼ To Delete an FCoE Port	59
▼ To Report Status on FCoE Ports	59
8 Configuring SAS Domains	61
SAS Multipathing Considerations	61
Dynamic Discovery of SAS Devices	61
Configuring SAS Boot Devices	62
9 Configuring IPFC SAN Devices	63
IPFC Considerations	63
Determining Fibre Channel Adapter Port Instances	64
Invoking and Configuring IPFC	66
▼ To Start a Network Interface Manually	66
▼ To Configure the Host for Automatic Plumbing Upon Reboot	67
10 Sun Solaris Express Developer's Edition OS x86/x64 Fibre Channel Operating System Booting Procedure	69
Setup Requirements	70
Installation Overview	70
Installation Procedure	71
▼ To Summarize the Procedure	71
▼ To Perform a CD-ROM, DVD ROM, or Network-Based OS Installation	71
Post-Installation Patches for the Sun Solaris Express Developer's Edition OS	78
11 Persistent Binding for Tape Devices	79
Persistent Binding Overview	79
Creating Tape Links	80

▼ Configuration Steps	81
A Manual Configuration for Fabric-Connected Devices	83
Manual Configuration	83
▼ To Configure Manually	84
Fabric Device Node Configuration	84
Ensuring That LUN Level Information Is Visible	85
Configuring Device Nodes Without Multipathing Enabled	86
▼ To Configure an Individual Device Without Multipathing	86
▼ To Configure Multiple Devices Without Multipathing	88
Configuring Device Nodes With Solaris Multipathing Software Enabled	89
▼ To Configure Individual Devices With Sun StorageTek Traffic Manager Software	90
▼ To Configure Multiple Devices With Multipathing	91
Unconfiguring Fabric Devices	93
Unconfiguring a Fabric Device	93
B Supported FC-HBA API	99
Supported Fibre Channel HBA API	99
C Multipathing Troubleshooting	103
System Fails to Boot During <code>stmsboot</code>	103
▼ To Recover From Single User Mode	103
System Crashes During <code>stmsboot</code>	104
▼ To Recover from a System Crash	104
Index	107

Preface

The *Solaris Express SAN Configuration and Multipathing Guide* provides an overview of the Sun StorageTek™ Traffic Manager software as an integrated part of the Sun Solaris™ Express Developer's Edition Operating System (OS). This guide includes instructions about how to install and configure the software and devices.

This guide is intended for system, storage and network administrators who create and maintain Fibre Channel (FC) storage area networks (SANs) and serial attached SCSI (SAS) Domains. A high level of expertise in the management and maintenance of SANS and SAS Domains is assumed.

Before You Read This Document

Before you read this book, read the latest Solaris Express Developer's Edition Operating System Release Notes.

How This Document Is Organized

[Chapter 1, “Product Overview,”](#) provides an overview of the Sun StorageTek Traffic Manager software.

[Chapter 2, “Fibre Channel Configuration Overview,”](#) provides an overview of configuring the Sun StorageTek Traffic Manager software.

[Chapter 3, “Administering Multipathing Devices Through `mpathadm` Commands,”](#) explains how to use the `mpathadm` command to enable multipathing administration through the ANSI standard Multipath Management API.

[Chapter 4, “Configuring Multipathing Software”](#) explains how to configure the multipathing feature for FC devices provided in the Solaris Express Developer's Edition OS.

[Chapter 5, “Configuring SAN Devices,”](#) provides high level overviews of the steps used to configure SAN devices.

[Chapter 6, “Configuring Virtual Fibre Channel Ports,”](#) provides the steps used to configure N Port ID Virtualization (NPIV) ports, also known as virtual Fibre Channel ports.

Chapter 7, “Configuring FCoE Ports,” provides the steps used to configure FCoE ports hosted on normal Ethernet interfaces.

Chapter 8, “Configuring SAS Domains,” provides high-level overviews of the steps used to configure SAS domains.

Chapter 9, “Configuring IPFC SAN Devices,” explains considerations when configuring IPFC SAN Devices.

Chapter 10, “Sun Solaris Express Developer's Edition OS x86/x64 Fibre Channel Operating System Booting Procedure” describes how to manually install the Solaris Express Developer's Edition OS, x86/x64 FC Operating System.

Chapter 11, “Persistent Binding for Tape Devices” describes how you can specify tape binding in the Solaris Express Developer's Edition OS while retaining the benefits of automatic discovery for disk-based devices.

Appendix A, “Manual Configuration for Fabric-Connected Devices” explains Manual configuration for Fabric Connected Devices.

Appendix B, “Supported FC-HBA API” provides details on Supported FC-HBA API.

Appendix C, “Multipathing Troubleshooting” provides trouble shooting information for potential problems that might occur while you run the Sun StorageTek Traffic Manager software.

Using UNIX Commands

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

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

Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

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

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- [Documentation](http://www.sun.com/documentation/) (<http://www.sun.com/documentation/>)
- [Support](http://www.sun.com/support/) (<http://www.sun.com/support/>)
- [Training](http://www.sun.com/training/) (<http://www.sun.com/training/>)

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. To share your comments, go to <http://docs.sun.com> and click Feedback.

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	machine_name%
C shell for superuser	machine_name#
Bourne shell and Korn shell	\$
Bourne shell and Korn shell for superuser	#

Product Overview

This chapter provides an overview of the Sun StorageTek Traffic Manager software, and covers the following topics:

- “About Sun StorageTek Traffic Manager” on page 11
- “What’s New?” on page 12
- “Software Features” on page 12
- “Supported Standards” on page 15

About Sun StorageTek Traffic Manager

The Sun StorageTek Traffic Manager software enables multipathing connections for Sun Solaris Express Developer's Edition OS hosts. Multipathing provides higher availability for storage devices through the use of multiple access paths. Fibre Channel (FC), Serial Attached SCSI (SAS), and Internet Small Computer Storage Interface (iSCSI) protocols are supported.

The Sun StorageTek Traffic Manager software resides on the host and identifies the storage devices on your SAN or SAS Domain. The software enables you to attach Fibre Channel storage in either loop, fabric, or point-to-point mode. The software provides a common interface for managing both Fibre Channel and SAS storage.

For information on how to configure targets and initiators in preparation for multipathing, see the following:

- *COMSTAR Administration Guide*, to configure a COMSTAR Fibre Channel target:
<http://opensolaris.org/os/project/comstar/Documentation/AdminGuide>
- Chapter 14 of the Solaris 10 *System Administration: Devices and File Systems Guide*, to configure Solaris iSCSI targets and initiators

What's New?

- Support for NPIV is included in the Solaris Express Community Edition (SXCE) Build 92, and all subsequent builds.

NPIV is a Fibre Channel facility that enables one Fibre Channel adapter to have many N Port IDs. Each N Port has a unique identity (port WWN and node WWN) on the SAN and can be used for zoning and LUN masking. Soft zoning, which you can use to group ports together by port WWN, is the preferred method of zoning. For more information, see [Chapter 6, “Configuring Virtual Fibre Channel Ports.”](#)

- Support for Fibre Channel over Ethernet (FCoE) is included in the SXCE Build 122, and all subsequent builds.

Fibre Channel over Ethernet (FCoE) is a proposed standard that is being developed by INCITS T11. The FCoE protocol specification maps Fibre Channel natively over Ethernet and is independent of the Ethernet forwarding scheme. The protocol provides I/O consolidation by preserving all Fibre Channel constructs, maintaining the same latency, security, and traffic management attributes of FC, while preserving investments in FC tools, training, and SANs. For more information, see [Chapter 7, “Configuring FCoE Ports.”](#)

Software Features

The Sun StorageTek Traffic Manager software enables multipathing connections for the Sun Solaris Express Developer's Edition OS hosts, with the following software features:

- [“Fibre Channel Software Features” on page 12](#)
- [“SAS Software Features” on page 13](#)
- [“Multipathing Software Features” on page 14](#)

Fibre Channel Software Features

The Sun StorageTek Traffic Manager software provides the following key features:

- **Dynamic Storage Discovery** – The software automatically recognizes devices and any modifications made to device configurations. This feature makes devices available to the system without requiring you to reboot or manually change information in configuration files.
- **Persistent Device Naming** – Devices that are configured within the software maintain their device naming through reboots or reconfiguration. The only exception to this policy are tape devices found in `/dev/rmt` that will not change unless they are removed, and then regenerated at a later date.
- **Fibre Channel Arbitrated Loop (FCAL) Support** – OpenBoot™ PROM (OBP) commands that are used on servers can access FCAL attached storage for scanning the FC loop.

- Fabric Booting – The Solaris OS supports booting from fabric devices as well as non-fabric Fibre Channel devices. Fabric topologies with Fibre Channel switches provide higher speed, more connections, and port isolation.
- FC-HBA Library – What was previously known as the Storage Networking Industry Association Fibre Channel host bus adapter (SNIA FC-HBA) library is now known as the FC-HBA library. The FC-HBA library application programming interface (API) enables management of FC HBAs and provides a standards-based interface for other applications (such as Sun StorEdge™ Enterprise Storage Manager) that can be used to gather information about FC HBAs.

Man pages for common FC-HBA APIs are included in the Solaris Express Developer's Edition OS in `libhbaapi(3LIB)`. For additional information about FC specifications, refer to <http://www.t11.org>.

- Fibre Channel Virtualization – N Port ID Virtualization (NPIV) is an extension to the Fibre Channel standard, which enables one Fibre Channel port to simulate many ports on the SAN. This is useful for virtualization environments such as xVM.
- Fibre Channel over Ethernet (FCoE) – A new T11 standard to transport encapsulated Fibre Channel frames over Enhanced Ethernet is now available. Solaris FCoE is a software implementation that is designed to work with normal Ethernet controllers.

SAS Software Features

- Dynamic Storage Discovery – The Solaris OS Storage Multipathing software automatically recognizes devices and any modifications made to device configurations. This makes devices available to the system without requiring you to reboot or manually change information in configuration files.
- Persistent Device Naming – Dynamic Storage Discovery Devices that are configured within the Solaris OS Storage Multipathing software maintain their device naming through reboots or reconfiguration.

Multipathing Software Features

The Solaris Express Developer's Edition OS contains the Sun StorageTek Traffic Manager software. The software is enabled by default for x86-based platforms and optional for SPARC-based platforms. The software contains the following features:

- **Path management** – The Sun StorageTek Traffic Manager software dynamically manages the paths to any storage devices it supports. The addition or removal of paths to a device is done automatically when a path is brought online or removed from a service. This feature enables hosts configured with the software to begin with a single path to a device and add more host controllers, increasing bandwidth and RAS, without changing device names or modifying applications. Sun storage does not require configuration files to manage or databases to keep current. For storage other than Sun, contact your storage vendor for methods to enable support and to ensure it is qualified with the Sun StorageTek Traffic Manager software.
- **Single Device Instances** – Unlike other multipathing solutions, the Sun StorageTek Traffic Manager software is fully integrated with the Solaris Express Developer's Edition OS. The software displays multipath devices as single device instances instead of as one device, or device link, per path. This feature reduces the cost of managing complex storage architectures utilities such as `format(1M)` or Sun Solaris OS Volume Manager, to “see” one representation of a storage device instead of a separate device for each path.
- **Failover support** – Implementing higher levels of RAS requires redundant host connectivity to storage devices. The Sun StorageTek Traffic Manager software manages the failure of storage paths while maintaining host I/O connectivity through available secondary paths.
- **Symmetrical/Asymmetrical device support**– The Sun StorageTek Traffic Manager software supports the following disk storage devices:
 - All Sun disk storage, both symmetric and asymmetric
 - All T10/T11 standards-compliant non-Sun symmetrical disk devices
 - Many non-Sun asymmetrical disk arrays

Note – T10 ALUA support has been added for asymmetrical devices that support this T10 standard. Contact your storage vendor to see whether your device is supported.

- **I/O load balancing** – In addition to providing simple failover support, the Sun StorageTek Traffic Manager software can use any active paths to a storage device to send and receive I/O. With I/O routed through multiple host connections, bandwidth can be increased by the addition of host controllers. The software uses a round-robin load-balancing algorithm by which individual I/O requests are routed to active host controllers in a series one after the other.

- **Queue depth** – SCSI storage arrays present storage to a host in the form of a logical unit number (LUN). LUNs have a finite set of resources available, such as the amount of data that can be stored, as well as the number of active commands that a device or LUN can process at one time. The number of active commands that can be issued before a device blocks further I/O is known as queue depth. When the Sun StorageTek Traffic Manager software is enabled, a single queue is created for each LUN regardless of the number of distinct or separate paths it may have to the host. This feature enables the disk driver to maintain and balance one queue to the LUN, effectively managing queue depth. No other multipathing software available for the Solaris Express Developer's Edition OS has this ability.
- **stmsboot command** – The Solaris Express Developer's Edition OS includes the `stmsboot (1M)` command, which allows you to enable or disable the Solaris OS Multipathing software on your boot device after the OS installation has been completed. This command is available for both SPARC and x86/x64 systems, and provides support for SAS multipathing with the Solaris Express Developer's Edition OS.

SAS multipathing is not enabled by default in the Solaris Express Developer's Edition OS for either SPARC or x86/x64 hosts, so use of the `stmsboot` command is required as a post-installation step if you wish to enable Multipathing.

SPARC hosts do not enable multipathing for Fibre Channel devices by default, so use of the `stmsboot` command is required as a post-installation step.

X86/x64 hosts do enable multipathing of Fibre Channel devices by default, so the use of the `stmsboot` command is not mandatory as a post-installation step.
- **Dynamic reconfiguration** – The Sun StorageTek Traffic Manager software supports the Solaris Express Developer's Edition OS dynamic reconfiguration (DR).

Supported Standards

The Sun StorageTek Traffic Manager software is based on open standards for communicating with devices and device management, ensuring interoperability with other standards-based devices and software. The following standards are supported by the Sun StorageTek Traffic Manager software:

- ANSI Standard: Information Technology – SNIA Multipath Management API Specification (ANSI INCITS 412-2006)
- T10 standards, including SCSI-3, SAM, FCP, SPC, and SBC
- T11.3 FC standards, including FC-PH, FC-AL, FC-LS, and FC-GS
- T11.5 storage management standards, including FC-HBA
- IETF standards, including RFC 2625
- Serial Attached SCSI-2 (SAS2)

Fibre Channel Configuration Overview

This chapter provides information about configuring the Solaris OS Fibre Channel (FC) and Storage Multipathing software. The following topics are discussed:

- “[Configuration Task Overview](#)” on page 17
- “[Configuration Considerations](#)” on page 18

Configuration Task Overview

This section provides a high level overview of the tasks that you should perform to configure the Sun StorageTek Traffic Manager software.

The following table lists the configuration tasks, description about each task and the reference information where you can find the procedures to perform each task.

TABLE 2-1 Tasks to Configure Solaris OS FC and Sun StorageTek Traffic Manager Software

Configuration Task	Task Description	Reference Information
Enable Multipathing	<p>Multipathing capability is provided by the Sun StorageTek Traffic Manager software in the Solaris Express Developer Edition OS. This software enables multipathing:</p> <ul style="list-style-type: none"> ▪ By default, for Fibre Channel-attached devices on x86/x64 platforms ▪ By manual configuration on SPARC platforms ▪ By manual configuration for SAS devices 	<p>The steps for enabling multipathing are outlined in Chapter 4, “Configuring Multipathing Software.”</p>

TABLE 2-1 Tasks to Configure Solaris OS FC and Sun StorageTek Traffic Manager Software
(Continued)

Configuration Task	Task Description	Reference Information
Configure Fibre Channel Devices	In the Solaris Express Developer Edition OS, FCAL, fabric, and point-to-point connected devices are made available to the host. This feature differs from previous versions of the Sun StorageTek Traffic Manager software running on the Solaris 8 or 9 OS. In those versions, manual configuration steps were required to achieve fabric-connected device availability on the Solaris host.	Chapter 5, “Configuring SAN Devices”
Configure Virtual Fibre Channel Ports	You can configure N Port ID Virtualization (NPIV) ports, which is especially useful for virtualization environments such as xVM.	Chapter 6, “Configuring Virtual Fibre Channel Ports”
Set up FC Boot Device	The Sun StorageTek Traffic Manager software enables a Solaris server to boot from a fabric device.	“SAN Device Considerations” on page 49
Configure IPFC SAN devices	You can configure host recognition of IPFC devices and implementation of IP over FC in a SAN. Configuration of IPFC depends on the instance of the Fabric Port (FP) driver for the FC adapter ports.	Chapter 9, “Configuring IPFC SAN Devices”
Configure SAS Devices	In the Solaris Express Developer Edition OS, multipathing of SAS devices is supported.	

Configuration Considerations

Before you configure the Sun StorageTek Traffic Manager software, consider the following:

- Configure ports and zones according to the vendor-specific documentation for storage and switches.
- LUN masking enables specific LUNs to be seen by specific hosts. See your vendor-specific storage documentation that describes masking.
- Power management needs to be disabled for hosts and devices on a SAN. See `power.conf` (4) for details about power management.
- The STMS Boot utility is included with the Sun StorageTek Traffic Manager software that manages the enable and disable process needed for SAN booting. Issuing the `smsboot` command allows the automatic update of `/etc/vfstab` and the dump configuration to reflect device name changes when enabling or disabling the Sun StorageTek Traffic Manager software. Note that the software is disabled by default with SPARC devices and enabled by default with x86 devices.

- For Sun Solaris Express Developer's Edition OS, fabric-connected devices are configured and made available to the host automatically during install and boot time.

Note – If you are performing an upgrade and want to make any Fibre Channel (FC) devices unavailable after upgrade, those devices have to be manually unconfigured using `cfgadm -c unconfigure` commands. However, to make those devices permanently unavailable to the system, you may want to consider using switch zoning or LUN masking. The changes made by `cfgadm -c unconfigure` do not persist after a reboot unless manual configuration for fabric-connected devices has been turned on. To find out how to disable the fabric devices discovery during boot or install, please refer to [Appendix A, “Manual Configuration for Fabric-Connected Devices”](#)

Administering Multipathing Devices Through `mpathadm` Commands

You determine and configure Solaris OS multipathing support by using `mpathadm` commands. In the Solaris Express Developer Edition OS, the `mpathadm` command enables multipathing administration through the ANSI standard Multipath Management API. The terms used in this chapter to denote a path, initiator port, target port, and logical unit are consistent with the T10 specification.

Administering Multipathing Devices

Perform the following tasks to administer any multipathing devices.

▼ To Get a List of Multipathing Support That Can Be Administered by the `mpathadm` Command

The multipathing support is identified with the Multipath Management API plug-in library name.

- List the multipathing support by typing the `mpathadm list mpath-support` command.

```
# mpathadm list mpath-support  
mpath-support: libmpscsi_vhci.so
```

(Only one result is shown in the preceding example.)

▼ To Get the Multipathing Support Properties

The multipathing support properties are associated with Multipath Management API plug-in properties.

- 1 List the multipathing support by typing the `mpathadm list mpath-support` command.

```
# mpathadm list mpath-support
mpath-support: libmpscsi_vhci.so
```

(Only one mpath support name result is shown in the preceding output list example.)

- 2 Choose the mpath support name from the output list.

- 3 List the multipathing support properties for a specified mpath support name by typing the `mpathadm show mpath-support mpath-support-name` command.

```
# mpathadm show mpath-support libmpscsi_vhci.so
mpath-support: libmpscsi_vhci.so
  Vendor: Sun Microsystems
  Driver Name: scsi_vhci
  Default Load Balance: round-robin
  Supported Load Balance Types:
    round-robin
    logical-block
  Allows To Activate Target Port Group Access: yes
  Allows path Override: no
  Supported Auto Failback Config: yes
  Auto Failback: on
  Failback Polling Rate (current/max): 0/0
  Supported Auto Probing config: no
  Auto Probing: NA
  probing Polling Rate (current/max): NA/NA
  Supported Devices:
    Vendor: SUN
    Product: T300
    Revision:
    Supported Load Balance Types:
      round-robin
    Vendor: SUN
    Product: T4
    Revision:
    Supported Load Balance Types:
      round-robin
```

The command output also shows a list of device products that is supported by the multipathing support software. The multipathing support `libmpscsi_vhci.so` supports T10 target port group compliant devices by default.

▼ To Get a List of Initiator Ports Discovered on the System

An initiator port is part of a path that consists of an initiator port, target port, and a logical unit.

- Run the `mpathadm list initiator-port` command.

```
# mpathadm list initiator-port
Initiator-Port: iqn.1986-03.com.sun:01:080020b7ac2b.437a3b3e,4000002affff
Initiator-Port: 2000000173018713
```

▼ To Get the Properties of a Particular Initiator Port

Display the properties of an initiator port by typing the show command.

- 1 Run the `mpathadm list initiator-port` command.

```
# mpathadm list initiator-port
Initiator-Port: iqn.1986-03.com.sun:01:080020b7ac2b.437a3b3e,4000002affff
Initiator-Port: 2000000173018713
```

- 2 Select an initiator port to use from the output.

- 3 Run the `mpathadm show initiator-port initiator-port-name` command.

```
# mpathadm show initiator-port 2000000173018713
initiator-port: 2000000173018713
  Transport Type: Fibre Channel
  OS Device File: /devices/pci@1f,2000/SUNW,glc@4/fp@0,0
```

▼ To Get a List of Logical Units Discovered on the System

The `mpathadm list lu` command displays the list of logical units attached to the initiator. This list displays the number of paths and the operational path count to each logical unit listed.

- Run the `mpathadm list lu` command.

```
# mpathadm list lu
/dev/rdisk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
```

```

Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
Total Path Count: 2
Operational Path Count: 2

```

▼ To Get the Properties and Associated Path Information of a Particular Logical Unit by Name

Display the list of logical units, along with the properties for each logical unit through the `mpathadm` command. The list of logical units contains logical unit names that in turn can be used to display the properties of a particular logical unit.

1 Run the `mpathadm list lu` command.

```

# mpathadm list lu
/dev/rdsk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
  Total Path Count: 2
  Operational Path Count: 2

```

2 Select a logical unit name from the output.

3 Run the `mpathadm show lu logical-unit-name` command.

```

# mpathadm show lu /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  mpath-support libmpscsi_vhci.so
  Vendor: SUN
  Product: T4
  Revision: 0301
  Name Type: SCSI Inquiry VPD page 83 type 3
  Name: 60003ba27d2120004204ac2b000dab00
  Asymmetric: yes
  Current Load Balance: round-robin
  Logical Unit Group ID: NA
  Auto Failback: on
  Auto Probing: NA

  Paths:
    Initiator Port Name: 2000000173818713

```



```

Target Port Name: 20030003ba27d212
Override Path: NA
Path State: OK
Disabled: no

```

```

Initiator Port Name: 2000000173018713
Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no

```

Target Port Groups:

```

ID: 2
Explicit Failover: yes
Access State: standby
Target Ports:
    Name: 20030003ba27d212
    Relative ID: 0

```

```

ID: 5
Explicit Failover: yes
Access State: active
Target Ports:
    Name: 20030003ba27d095
    Relative ID: 0

```

▼ To List All Logical Units Associated With a Particular Target Port

To list paths through a target port use the `show` command of a logical unit. First, retrieve the list of target ports associated with a logical unit. Then, use the target port to list all logical units associated with this target port.

1 Run the `mpathadm list lu` command.

```

# mpathadm list lu
/dev/rdisk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
    Total Path Count: 2
    Operational Path Count: 2

```

2 Select a logical unit name from the output.**3 Run the `mpathadm show lu` command.**

```
# mpathadm show lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
```

```
logical unit: /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
```

```
mpath-support libmpscsi_vhci.so
```

```
Vendor: SUN
```

```
Product: T4
```

```
Revision: 0301
```

```
Name Type: SCSI Inquiry VPD page 83 type 3
```

```
Name: 60003ba27d2120004204ac2b000dab00
```

```
Asymmetric: yes
```

```
Current Load Balance: round-robin
```

```
Logical Unit Group ID: NA
```

```
Auto Failback: on
```

```
Auto Probing: NA
```

```
Paths:
```

```
Initiator Port Name: 2000000173818713
```

```
Target Port Name: 20030003ba27d212
```

```
Override Path: NA
```

```
Path State: OK
```

```
Disabled: no
```

```
Initiator Port Name: 2000000173018713
```

```
Target Port Name: 20030003ba27d095
```

```
Override path: NA
```

```
Path State: OK
```

```
Disabled: no
```

```
Target Port Groups:
```

```
ID: 2
```

```
Explicit Failover: yes
```

```
Access State: standby
```

```
Target Ports:
```

```
    Name: 20030003ba27d212
```

```
    Relative ID: 0
```

```
ID: 5
```

```
Explicit Failover: yes
```

```
Access State: active
```

```
Target Ports:
```

```
    Name: 20030003ba27d095
```

```
    Relative ID: 0
```

4 Select a target port from the output list.

5 Run the `mpathadm list lu -t` command.

```
# mpathadm list lu -t 20030003ba27d212
/dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
    Total Path Count: 2
    Operational Path Count: 2
```

▼ To List a Logical Unit With a Particular Name Derived From the Hardware Used by the System**1 Run the `mpathadm list lu` command.**

```
# mpathadm list lu
/dev/rdsk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
    Total Path Count: 2
    Operational Path Count: 2
```

2 Select a logical unit name from the output.**3 Run the `mpathadm show lu logical-unit-name` command.**

```
# mpathadm show lu /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
mpath-support libmpscsi_vhci.so
Vendor: SUN
Product: T4
Revision: 0301
Name Type: SCSI Inquiry VPD page 83 type 3
Name: 60003ba27d2120004204ac2b000dab00
Asymmetric: yes
Current Load Balance: round-robin
Logical Unit Group ID: NA
Auto Failback: on
Auto Probing: NA

Paths:
```

```

Initiator Port Name: 2000000173818713
Target Port Name: 20030003ba27d212
Override Path: NA
Path State: OK
Disabled: no

```

```

Initiator Port Name: 2000000173018713
Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no

```

Target Port Groups:

```

ID: 2
Explicit Failover: yes
Access State: standby
Target Ports:
  Name: 20030003ba27d212
  Relative ID: 0

```

```

ID: 5
Explicit Failover: yes
Access State: active
Target Ports:
  Name: 20030003ba27d095
  Relative ID: 0

```

- 4 Select an *inquiry-string-name* from the output.
- 5 Run the `mpathadm list lu -n inquiry-string-name` command.

```

# mpathadm list lu -n 60003ba27d2120004204ac2b000dab00
mpath-support: libmpscsi_vhci.so
/dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
  Operational Path Count: 2

```

▼ To Configure Auto-Failback for a Particular Multipathing Support

Symmetrical devices allow for automatic failback to a possible optimal path. If a failover occurs on the initial path, the standby path becomes the new online path. Usually the standby path is a suboptimal path. When auto-failback is enabled, the initial path comes back online, failover to the initial path automatically occurs.

1 Run the `mpathadm list mpath-support` command.

```
# mpathadm list mpath-support
mpath-support: libmpscsi_vhci.so
```

2 Select an `mpath` support name from the output list.**3 Run the `mpathadm modify mpath-support -a [on | off] mpath-support-name` command.**

```
# mpathadm modify mpath-support -a off libmpscsi_vhci.so
```

4 Run the `mpathadm show mpath-support mpath-support-name` command to check the change.

```
# mpathadm show mpath-support libmpscsi_vhci.so
mpath-support: libmpscsi_vhci.so
Vendor: Sun Microsystems
Driver Name: scsi_vhci
Default Load Balance: round-robin
Supported Load Balance Types:
    round-robin
    logical-block
Allows To Activate Target Port Group Access: yes
Allows path Override: no
Supported Auto Failback Config: yes
Auto Failback: on
Failback Polling Rate (current/max): 0/0
Supported Auto Probing config: no
Auto Probing: NA
probing Polling Rate (current/max): NA/NA
Supported Devices:
Vendor: SUN
Product: T300
Revision:
Supported Load Balance Types:
    round-robin
Vendor: SUN
Product: T4
Revision:
Supported Load Balance Types:
    round-robin
```

Note – The auto-failback display initiated by the `mpathadm modify` command setting is effective while the system is running. However, to keep the changed setting persistent, you must update the `/kernel/drv/scsi_vhci.conf` file. Refer to [“Configuring Automatic Failback” on page 47](#).

▼ To Fail Over a Logical Unit

This operation is applicable only to devices in the following two categories:

- Asymmetric devices with a proprietary failover mechanism recognized and supported by multipathing support
- Devices conforming to the T10 Standard Target Port Group Support `libmpscsi_vhci.so` and providing *explicit* mode asymmetric logical unit access

1 Run the `mpathadm list lu` command.

```
# mpathadm list lu
/dev/rdisk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
  Total Path Count: 2
  Operational Path Count: 2
```

2 Select a logical unit name from the output.

3 Run the `mpathadm show lu logical-unit-name` command.

```
# mpathadm show lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  mpath-support libmpscsi_vhci.so
  Vendor: SUN
  Product: T4
  Revision: 0301
  Name Type: SCSI Inquiry VPD page 83 type 3
  Name: 60003ba27d2120004204ac2b000dab00
  Asymmetric: yes
  Current Load Balance: round-robin
  Logical Unit Group ID: NA
  Auto Failback: on
  Auto Probing: NA

  Paths:
    Initiator Port Name: 2000000173818713
    Target Port Name: 20030003ba27d212
    Override Path: NA
    Path State: OK
    Disabled: no

    Initiator Port Name: 2000000173018713
```

```

Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no

```

```

Target Port Groups:
ID: 2
Explicit Failover: yes
Access State: standby
Target Ports:
    Name: 20030003ba27d212
    Relative ID: 0

ID: 5
Explicit Failover: yes
Access State: active
Target Ports:
    Name: 20030003ba27d095
    Relative ID: 0

```

4 Run the `mpathadm failover lu logical-unit-name` command.

```
# mpathadm failover lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
```

If this operation is successful, the access state of the devices' target port groups changes as a result of the logical unit failover.

5 Run the `mpathadm show lu logical-unit-name` command to verify the access state change.

```

# mpathadm show lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  mpath-support libmpscsi_vhci.so
  Vendor: SUN
  Product: T4
  Revision: 0301
  Name Type: SCSI Inquiry VPD page 83 type 3
  Name: 60003ba27d2120004204ac2b000dab00
  Asymmetric: yes
  Current Load Balance: round-robin
  Logical Unit Group ID: NA
  Auto Failback: on
  Auto Probing: NA

  Paths:
    Initiator Port Name: 2000000173818713
    Target Port Name: 20030003ba27d212
    Override Path: NA
    Path State: OK
    Disabled: no

```

```
Initiator Port Name: 2000000173018713
Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no
```

Target Port Groups:

```
ID: 2
Explicit Failover: yes
Access State: active
Target Ports:
  Name: 20030003ba27d212
  Relative ID: 0
```

```
ID: 5
Explicit Failover: yes
Access State: standby
Target Ports:
  Name: 20030003ba27d095
  Relative ID: 0
```

▼ To Enable a Path

If the path to a logical unit is disabled, the `enable` command changes the path back to enabled. The full path must be specified using the initiator port name, target port name, and the logical unit. To verify the change, run the `show` command for the logical unit.

1 Run the `mpathadm list lu` command.

```
# mpathadm list lu
/dev/rdisk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
  Total Path Count: 2
  Operational Path Count: 2
```

2 Select a logical unit name from the output.

3 Run the `mpathadm show lu logical-unit-name` command.

```
# mpathadm show lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
```



```

mpath-support libmpscsi_vhci.so
Vendor: SUN
Product: T4
Revision: 0301
Name Type: SCSI Inquiry VPD page 83 type 3
Name: 60003ba27d2120004204ac2b000dab00
Asymmetric: yes
Current Load Balance: round-robin
Logical Unit Group ID: NA
Auto Failback: on
Auto Probing: NA

```

Paths:

```

Initiator Port Name: 2000000173818713
Target Port Name: 20030003ba27d212
Override Path: NA
Path State: OK
Disabled: no

```

```

Initiator Port Name: 2000000173018713
Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no

```

Target Port Groups:

```

ID: 2
Explicit Failover: yes
Access State: standby
Target Ports:
  Name: 20030003ba27d212
  Relative ID: 0

```

```

ID: 5
Explicit Failover: yes
Access State: active
Target Ports:
  Name: 20030003ba27d095
  Relative ID: 0

```

- 4 **Select an initiator port and a target port name.**
- 5 **Run the `mpathadm enable path -i initiator-port-name -t target-port-name -l logical-unit-name` command.**

```

# mpathadm enable path -i 2000000173018713 -t 20030003ba27d095 \
  -l /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2

```

▼ To Disable a Path

This operation makes the path unavailable for the logical unit, regardless of its operational state.

Note – The disabled state is not persistent across rebooting. If the path is operational before the next boot sequence, it is enabled by default. This operation is not allowed when the given path is the last operational path remaining.

1 Run the `mpathadm list lu` command.

```
# mpathadm list lu
/dev/rdsk/c4t60020F20000035AF4267CCCB0002CEE2d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdsk/c4t60003BA27D2120004204AC1A000F1D6Cd0s2
  Total Path Count: 2
  Operational Path Count: 2
```

2 Select a logical unit name from the output.

3 Run the `mpathadm show lu logical-unit-name` command.

```
# mpathadm show lu /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
logical unit: /dev/rdsk/c4t60003BA27D2120004204AC2B000DAB00d0s2
  mpath-support libmpscsi_vhci.so
  Vendor: SUN
  Product: T4
  Revision: 0301
  Name Type: SCSI Inquiry VPD page 83 type 3
  Name: 60003ba27d2120004204ac2b000dab00
  Asymmetric: yes
  Current Load Balance: round-robin
  Logical Unit Group ID: NA
  Auto Failback: on
  Auto Probing: NA

  Paths:
    Initiator Port Name: 2000000173818713
    Target Port Name: 20030003ba27d212
    Override Path: NA
    Path State: OK
    Disabled: no

    Initiator Port Name: 2000000173018713
```

```

Target Port Name: 20030003ba27d095
Override path: NA
Path State: OK
Disabled: no

```

Target Port Groups:

```

ID: 2
Explicit Failover: yes
Access State: standby
Target Ports:
  Name: 20030003ba27d212
  Relative ID: 0

```

```

ID: 5
Explicit Failover: yes
Access State: active
Target Ports:
  Name: 20030003ba27d095
  Relative ID: 0

```

- 4 **Select an initiator port and a target port name.**
- 5 **Run the `mpathadm disable path -i initiator-port-name -t target-port-name -l logical-unit-name` command.**

```

# mpathadm disable path -i 2000000173018713 -t 20030003ba27d095 \
-l /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2

```


Configuring Multipathing Software

This chapter explains how to configure the Sun StorageTek Traffic Manager software that is provided in the Solaris Express Developer Edition OS for Fibre Channel (FC) devices. It also provides the considerations while enabling or disabling the multipathing feature on SPARC based systems, x86 based systems, per-port basis and third-party storage devices.

The following topics are discussed:

- “Configuring Multipathing” on page 37
- “Multipathing Considerations” on page 38
- “Enabling and Disabling Multipathing” on page 38
- “Enabling or Disabling Multipathing on a Per-Port Basis” on page 41
- “Configuring Third-Party Storage Devices” on page 43
- “Configuring Automatic Failback” on page 47

Configuring Multipathing

The multipathing feature for FC devices can be configured to control all Sun supported FC HBAs. Multipathing is disabled by default for FC devices on SPARC based systems, but is enabled by default on x86 based systems. Configuration of the multipathing software depends on how you intend to use your system.

Note – The multipathing feature is not available for parallel SCSI devices but is available for FC disk devices. Multipathing is not supported on tape drives or libraries, or on IP over FC.

Before you configure third-party devices for multipathing, be aware of the following:

Multipathing Considerations

Before you change multipathing configuration note the following considerations. Then follow the instructions for your machine architecture (SPARC or x86) described in the subsequent sections. Some devices need to be properly configured to work with the multipathing software. Refer to your storage array documentation for details on the device specific configuration for your device.

- Device Specific Considerations
- Device Name Change Considerations

In the `/dev` and `/devices` trees, multipath-enabled devices receive new names that indicate that they are under multipath control. A device therefore will have a different name from its original name when it is under multipath control.

Device name with multipath disabled:

```
/dev/dsk/c1t1d0s0
```

Device name with multipath enabled:

```
/dev/dsk/c0t60003BA27D5170003E5D2A7A0007F3D2d0s0
```

Therefore, applications that use device names directly must be configured to use the new names whenever you change multipath configuration from disabled to enabled or vice-versa.

The system's `/etc/vfstab` file and the dump configuration also contain references to device names. On both SPARC and x86/x64-based systems the `stmsboot (1m)` command described in the following sections automatically updates the `/etc/vfstab` file dump configuration with the new device names. If you have application-dependent file systems which are not listed in the file `/etc/vfstab`, you can use the `stmsboot` command to determine the mapping between the old and new device paths.



Caution – if you have run `devfsadm -C` or performed a reconfiguration reboot then the old device paths will not exist and the `stmsboot -L` command will fail to provide this information.

Enabling and Disabling Multipathing

With the release of the Solaris Express Developer Edition OS, the `stmsboot (1m)` command is now available for use on SPARC, x86, and x64 based systems to enable or disable multipathing for Fibre Channel (FC) and SAS devices. The `stmsboot` command updates the `/etc/vfstab` file and dump configuration to reflect device name changes during the next reboot. You do not need to manually edit the `/kernel/drv/fp.conf` or `/kernel/drv/mpt.conf` files.

The following considerations apply to the `stmsboot -e`, `-d`, and `-u` options:

- You should reboot immediately after running the `stmsboot` command.
- Ensure that `eeprom(1m)` boot device is set to boot from your current boot device. This is required since the `stmsboot` command reboots the machine to complete the operation.
- The `stmsboot` command saves a copy of your original `/kernel/drv/fp.conf`, `/kernel/drv/mpt.conf` and `/etc/vfstab` files before modifying them. See [Appendix C, “Multipathing Troubleshooting,”](#) if you encounter unexpected problems while using `stmsboot`.

Note – The behavior of the `stmsboot` command in the Solaris Express Developer Edition OS differs from the `stmsboot` command provided by Sun StorEdge SAN Foundation software in the Solaris 8 and Solaris 9 OS. In the Solaris 8 and Solaris 9 OS, the `stmsboot` command is used to enable or disable multipathing on the boot device only for SPARC based hosts only. In the Solaris Express Developer Edition OS, the command is used to enable or disable multipathing on all capable devices.

The `stmsboot` command is available on both SPARC and x86/x64-based systems from the Solaris Express Developer Edition OS onwards. In the Solaris Express Developer Edition OS the `stmsboot` command is used to enable, disable or update multipathing for Serial-Attached SCSI devices as well as FC devices.

▼ To Enable Multipathing

Perform the following steps to enable multipathing on all multipath-capable devices on SPARC or x86/x64 based systems. If you want to enable multipathing only on specific FC or SAS HBA ports, refer to [“Enabling or Disabling Multipathing on a Per-Port Basis”](#) on page 41.

The multipathing software automatically recognizes Sun supported devices. If you want to enable multipathing on third-party devices, add entries to the `/kernel/drv/scsi_vhci.conf` file as described by *Device Name Change Considerations* in [“Multipathing Considerations”](#) on page 38.

1 Run the `stmsboot -e` command:

```
# stmsboot -e
```

```
WARNING: stmsboot operates on each supported multipath-capable controller
         detected in a host. In your system, these controllers are
```

```
/devices/pci@780/pci@0/pci@8/SUNW,qlc@0/fp@0,0
/devices/pci@780/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2
```

```
/devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
/devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
```

If you do NOT wish to operate on these controllers, please quit `stmsboot` and re-invoke with `-D { fp | mpt }` to specify which controllers you wish to modify your multipathing configuration for.

```
Do you wish to continue? [y/n] (default: y) y
Checking mpzio status for driver fp
Checking mpzio status for driver mpt
WARNING: This operation will require a reboot.
Do you want to continue ? [y/n] (default: y) y
The changes will come into effect after rebooting the system.
Reboot the system now ? [y/n] (default: y) y
```

Note – During the reboot, `/etc/vfstab` and the dump configuration are updated to reflect the device name changes.

- 2 (Optional) After the reboot, if necessary, configure your applications to use new device names as described by *Device Name Change Considerations* in [“Multipathing Considerations” on page 38](#).

▼ To Disable Multipathing

Perform the following steps to disable multipathing on all multipath-capable devices on SPARC or x86/x64 based systems. If you want to disable multipathing only on specific FC or SAS HBA ports, refer to [“Enabling or Disabling Multipathing on a Per-Port Basis” on page 41](#).

- 1 Run the `stmsboot -d` command.

```
# stmsboot -d
```

WARNING: `stmsboot` operates on each supported multipath-capable controller detected in a host. In your system, these controllers are

```
/devices/pci@780/pci@0/pci@8/SUNW,qlc@0/fp@0,0
/devices/pci@780/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1
/devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2
/devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
/devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0
```

If you do NOT wish to operate on these controllers, please quit `stmsboot` and re-invoke with `-D { fp | mpt }` to specify which controllers you wish to modify your multipathing configuration for.


```

Do you wish to continue? [y/n] (default: y) y
Checking mpzio status for driver fp
Checking mpzio status for driver mpt
WARNING: This operation will require a reboot.
Do you want to continue ? [y/n] (default: y) y
The changes will come into effect after rebooting the system.
Reboot the system now ? [y/n] (default: y) y

```

Note – During the reboot, `/etc/vfstab` and the dump configuration are updated to reflect the device name changes.

- 2 (Optional) After the reboot, if necessary, configure your applications to use new device names as described by *Device Name Change Considerations* in [“Multipathing Considerations” on page 38](#).

Enabling or Disabling Multipathing on a Per-Port Basis

Multipathing can be enabled or disabled on specific Fibre Channel Host Bus Adapter (HBA) controller ports. If you enable multipathing on a specific HBA port controller port, all supported devices connected to that controller port will be enabled for multipath operation.

The following procedure applies to both SPARC and x86 based machines.

Port Configuration Considerations

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

- FC global and per-port multipath settings are specified in the file `/kernel/drv/fp.conf`. Per-port multipath settings have priority over the global setting. This means that if global multipathing is enabled by a specific port has been disabled for multipathing, the port will not be available in the multipathing configuration. Conversely, even if global multipathing has been disabled, specific ports may be enabled for multipathing if they are listed in the appropriate `driver.conf(4)` file.
- Load balancing is controlled by the global `load-balance` property in `/kernel/drv/scsi_vhci.conf` file and is not controlled on a per-port basis.
- If a device has more than one path to the host, all paths to the device must be configured with multipathing enabled or disabled.
- Configuring multipathing by port enables the multipathing software to coexist with other multipathing solutions like Symantec (VERITAS) Dynamic Multipathing (DMP), or EMC PowerPath. However, devices and paths should not be shared between the multipathing software and other multipathing solutions.

▼ To Configure Multipathing by Port

Depending on how many ports you want the multipathing software to control, you can enable or disable multipathing globally or for specified ports.

- 1 **Log in as root (su - root).**
- 2 **Determine the HBA controller ports that you want the multipathing software to control.**

For example, to select the desired device, perform an `ls -l` command on `/dev/cfg` directory. The following example shows the `ls -l` command output.

```
lrwxrwxrwx 1 root  root  50 Jan 29 21:33 c0 ->
    ../../devices/pci@7c0/pci@0/pci@1/pci@0/ide@8:scsi
lrwxrwxrwx 1 root  root  61 Jan 29 21:33 c1 ->
    ../../devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@1:scsi
lrwxrwxrwx 1 root  root  61 Jan 29 21:33 c2 ->
    ../../devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2:scsi
lrwxrwxrwx 1 root  root  53 Jan 29 21:33 c3 ->
    ../../devices/pci@7c0/pci@0/pci@9/LSILogic,sas@0:scsi
lrwxrwxrwx 1 root  root  54 Apr 16 20:28 c5 ->
    ../../devices/pci@780/pci@0/pci@8/SUNW,qlc@0/fp@0,0:fc
lrwxrwxrwx 1 root  root  56 Apr 16 20:28 c6 ->
    ../../devices/pci@780/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0:fc
```

Note – Controllers c5 and c6 are ports A and B on a dual-port FC HBA. Controllers c1 and c3 are single port SAS HBA ports. Controller c2 is the internal SAS controller in a Sun Fire™ T2000 server.

Once you have determined the port, or ports, you want to explicitly enable or disable multipathing for, go to the next step.

- 3 **Enable or disable specific ports as follows:**

- **FC HBA ports**

- a. **Edit file** `/kernel/drv/fp.conf`

- b. **Add the following line for each FC HBA port that is to be enabled:**

```
name="fp" parent="parent name" port=port-number mpzio-disable="no";
```

where *parent name* is the port device name, and *port-number* is the FC HBA port number.

For example, the following entries disable multipathing on all FC HBA controller ports except for the two specified ports:

```
mpxio-disable="yes";
name="fp" parent="/pci@6,2000/SUNW,qlc@2" port=0 mpxio-disable="no";
name="fp" parent="/pci@13,2000/pci@2/SUNW,qlc@5" port=0 mpxio-disable="no";
```

c. Add the following line for each FC HBA port that is to be disabled:

```
name="fp" parent="parent name" port=port-number mpxio-disable="yes";
```

For example:

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

4 Type `stmsboot -u` to start the reboot and configuration process.

You are prompted to reboot. During the reboot, the `/etc/vfstab` file and your dump device configuration are updated to reflect any device name changes.

5 (Optional) After the reboot, if necessary, configure your applications to use new device names as described in “Multipathing Considerations” on page 38.

Configuring Third-Party Storage Devices

Note – Before configuring any third-party device, ensure that they are supported. Refer to your third-party user documentation, or third-party vendor for information on proper vendor and product IDs, modes and various settings required for the device to work with multipathing software.

Third-Party Device Configuration Considerations

Before you configure third-party devices for multipathing, be aware of the following:

- The device must support the `REPORT_LUNS` SCSI command, and `SCSI-3 INQUIRY` command VPD Device Identification Page (0x83).
- You will need the vendor ID (VID) and product ID (PID) of the device. You can obtain them by using the `format` command followed by the `inquiry` option on your system. See the `format(1M)` man page.

When multipathing is enabled, the multipath access still depends on a device specific `scsi_vhci` failover implementation accepting the device. The default way this occurs, automatically, is for the `scsi_vhci` code to call a *probe* function in each failover implementation, looking for the first *probe* result that indicates the device is supported.

A *probe* implementation determines support based on some combination of `scsi_inquiry(9S)` data. A device with INQUIRY data indicating T10 Target-Port-Group-Support (TPGS) compliance will use the standards-based TPGS failover implementation. For noncompliant devices, a failover implementation's probe will typically determine support based on VID/PID match against a private compiled-in table.

To override the *probe* process, the `scsi_vhci.conf` file supports a `scsi-vhci-failover-override` property. The value of `scsi-vhci-failover-override` can be used to establish support for a device not currently accepted by *probe*, override *probe* support, or disable multipath support for a device.

Configuring Third-Party Storage Devices: New Devices

Multipathing can be configured on third-party symmetric storage devices. A symmetric storage device is defined as a storage device in which all paths to the storage device are active and I/O commands can be issued through any path.

▼ To Configure New Third-Party Devices

Perform the following steps to configure third-party devices if your system already has multipathing enabled. If your system has multipathing disabled, you can configure third-party devices while enabling multipathing as described earlier in this chapter.

1 Add the vendor ID and product ID entries to the `/kernel/drv/scsi_vhci.conf` file.

The vendor ID and product ID are the vendor and product identification strings that the device returns in SCSI INQUIRY data. The vendor ID must be eight characters long. You must specify all eight characters even if the trailing characters are spaces.

The product ID can be up to 16 characters long.

```
scsi-vhci-failover-override =
"VendorID1ProductID1", "f_sym",
"VendorID2ProductID2", "f_sym",
...
"VendorIDnProductIDn", "f_sym";
```

Note that the entries are separated by the `;` character and the last vendor/product entry is terminated by the `;` character.

For example, to add a device from a “ACME” vendor with a product ID of “MSU”, and to add a device from vendor “XYZ” vendor with a product ID of “ABC”, you would add the following lines to file `/kernel/drv/scsi_vhci.conf`:

To add a device from a “XYZ” vendor with a product ID of “ABC”, you would add:

```
scsi-vhci-failover-override =
"ACME    MSU", "f_sym",
"XYZ     ABC", "f_sym";
```

- 2 **Save and exit the `scsi_vhci.conf` file.**
- 3 **Type `stmsboot -u` to start the reboot and configuration process.**
You are prompted to reboot. During the reboot, the `/etc/vfstab` file and the dump configuration are updated to reflect the device name changes.
- 4 **If necessary, perform device name updates as described in [“Enabling and Disabling Multipathing” on page 38](#).**

Configuring Third-Party Storage Devices: Disabling Devices

Multipathing can be disabled for all devices of a certain Vendor ID/Product ID combination from Inquiry data. This exclusion is specified in the `scsi_vhci.conf` file.

▼ To Disable Third-Party Devices

- 1 **Add the vendor ID and product ID entries to the `/kernel/drv/scsi_vhci.conf` file.**

The vendor ID and product ID are the vendor and product identification strings that the device returns in SCSI INQUIRY data. The vendor ID must be eight characters long. You must specify all eight characters even if the trailing characters are spaces. The product ID can be up to 16 characters long.

```
scsi-vhci-failover-override =
"VendorID1ProductID1", "NONE",
"VendorID2ProductID2", "NONE",
...
"VendorIDnProductIDn", "NONE";
```

The entries in the preceding example are separated by the `'`, `'` character and the last vendor/product entry is terminated by the `'`, `'` character. For example, to add a device from vendor “ACME” with a product ID of “MSU,” and a vendor device from vendor “XYZ” with product ID “ABC,” you would add the following lines to the file `/kernel/drv/scsi_vhci.conf`:

```
scsi-vhci-failover-override =
"ACME    MSU", "NONE",
"XYZ     ABC", "NONE";
```

- 2 **Save and exit the `scsi_vhci.conf` file.**
- 3 **Type `stmsboot -u` to start the reboot and configuration process.**
You are prompted to reboot. During the reboot, the `/etc/vfstab` file and the dump configuration are updated to reflect the device name changes.

- 4 If necessary, perform the device name updates as described in [“Enabling and Disabling Multipathing” on page 38](#).

Displaying Device Name Changes

The `stmsboot(1M)` command, available on Solaris Express Developer's Edition OS SPARC, x86 and x64 platforms, can be used to display the mapping between non-multipathed and multipathed device names, after the changes made to the multipath configuration by invoking `stmsboot` with the `-e`, `-d`, or `-u` option. Both the non-multipathed and the multipathed device names must exist in order to show the mapping.

The `-L` option displays the mapping of devices on all controllers.

```
# stmsboot -L
non-STMS device name          STMS device name
-----
/dev/rdisk/c2t8d0            /dev/rdisk/c10t500000E01046DEE0d0
/dev/rdisk/c2t0d0            /dev/rdisk/c10t500000E01046B070d0
/dev/rdisk/c2t3d0            /dev/rdisk/c10t20000020372A40AFd0
/dev/rdisk/c2t12d0           /dev/rdisk/c10t500000E01046DEF0d0
/dev/rdisk/c2t11d0           /dev/rdisk/c10t500000E01046E390d0
/dev/rdisk/c3t8d0            /dev/rdisk/c10t500000E01046DEE0d0
/dev/rdisk/c3t0d0            /dev/rdisk/c10t500000E01046B070d0
/dev/rdisk/c3t3d0            /dev/rdisk/c10t20000020372A40AFd0
/dev/rdisk/c3t12d0           /dev/rdisk/c10t500000E01046DEF0d0
/dev/rdisk/c3t11d0           /dev/rdisk/c10t500000E01046E390d0
```

The `-l` option displays the mapping of devices on only the specified controller. The following example displays the mapping of controller 3.

```
# stmsboot -l3
non-STMS device name          STMS device name
-----
/dev/rdisk/c3t8d0            /dev/rdisk/c10t500000E01046DEE0d0
/dev/rdisk/c3t0d0            /dev/rdisk/c10t500000E01046B070d0
/dev/rdisk/c3t3d0            /dev/rdisk/c10t20000020372A40AFd0
/dev/rdisk/c3t12d0           /dev/rdisk/c10t500000E01046DEF0d0
/dev/rdisk/c3t11d0           /dev/rdisk/c10t500000E01046E390d0
```

Configuring Automatic Failback

Some storage devices have controllers configured as PRIMARY and SECONDARY as part of the array configuration. The secondary paths may operate at a lower performance level than the primary paths. The multipathing software uses the primary path to talk to the storage device and keeps the secondary path on standby. In the event of a primary path failure, the multipathing software automatically directs all I/O traffic over the secondary path, with the primary path taken offline. This process is called a “failover” operation. When the failure associated with the primary path has been repaired, the multipathing software automatically directs all I/O traffic over the primary path and keeps the secondary path standby as before. This process is called a “failback” operation.

You can disable the automatic failback operation so the multipathing software does not automatically failback to the primary path. Later, after the failure associated with the primary path has been repaired, you can do a manual failback operation using the `luxadm(1M)` command.

▼ To Configure Automatic Failback

- 1 **Open the `/kernel/drv/scsi_vhci.conf` file in a text editor.**

- 2 **To disable automatic failback capability, change the `auto-failback` entry to:**

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

To enable automatic failback capability (if you disabled it before), change the `auto-failback` entry to:

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

- 3 **Save and exit the file.**

- 4 **Reboot the system.**

```
# shutdown -g0 -y -i6
```


Configuring SAN Devices

This chapter provides the high level overview of the steps used to configure SAN devices. In the Solaris Express Developer Edition OS, FCAL, fabric, and point-to-point connected devices are made available to the host automatically. This feature differs from previous versions of the SAN Foundation software running on the Solaris 8 or the Solaris 9 OS. In those versions, manual configuration steps were required to achieve fabric-connected device availability on the host. To find out how to manually configure fabric-connected devices, refer to [Appendix A, “Manual Configuration for Fabric-Connected Devices.”](#)

The following topics are discussed:

- “SAN Device Considerations” on page 49
- “Adding SAN Devices” on page 50
- “Configuring Fabric Boot Devices on SPARC” on page 51

SAN Device Considerations

Be aware of the following considerations before configuring the Sun StorageTek Traffic Manager software.

- Configure ports and zones according to the vendor-specific documentation for storage and switches.
- With the Solaris Express Developer Edition OS you no longer have to manually configure devices unless you enable manual device configuration.
- LUN masking enables specific LUNs to be seen by specific hosts. See your vendor-specific storage documentation that describes masking.
- Turn off power management on servers connected to the SAN to prevent unexpected results as one server attempts to power down a device while another attempts to gain access. See `power.conf(1M)` man page for details about power management.

- Connect arrays and other storage devices to the SAN with or without multipathing capability. The Sun StorageTek Traffic Manager software is an associated application that is bundled with the product name.
- The STMS Boot utility is included with the Sun StorageTek Traffic Manager software that manages the enable and disable process needed for SAN booting. Issuing the `stmsboot` command allows the automatic update of `/etc/vfstab` and the dump configuration to reflect device name changes when enabling or disabling the Sun StorageTek Traffic Manager software. Note that the software is disabled by default with SPARC devices and enabled by default with x86 devices.

Adding SAN Devices

Adding and removing SAN devices requires knowledge of the following commands:

- `luxadm(1M)`
- `format(1M)`
- `fscck(1M)`
- `newfs(1M)`
- `cfgadm(1M)` and `cfgadm_fp(1M)`

Note – If you use the `format` command when the Sun StorageTek Traffic Manager software multipathing feature is enabled, you will see only one instance of a device identifier for each LUN. Without the Sun StorageTek Traffic Manager software enabled, you will see one identifier for each path.

The `cfgadm` and `cfgadm_fp` commands are used most frequently to configure storage devices on a SAN. Refer to the appropriate man page for detailed instructions about how to use each command.

▼ To Add a SAN Device

- 1 Create the LUNs desired on the SAN device.
- 2 If necessary, apply LUN masking for HBA control on the SAN device.
- 3 Connect the storage device to the system.
- 4 If necessary, create port-based or WWN zones on the switch on the SAN device.
- 5 Run the `fscck` or `newfs` commands on the device, if used for file systems.

6 Mount any existing file systems available on the storage device's LUNs or disk groups.

Note – You might need to run the `fsck` command to repair any errors in the LUNs listed in the `/etc/vfstab` file.

Configuring Fabric Boot Devices on SPARC

The Sun StorageTek Traffic Manager software allows a Solaris Express Developer Edition OS SPARC server to be booted from a fabric disk device.

Boot Considerations

Fabric boot devices can be added through the Solaris Express Developer Edition OS interactive installation GUI or text installer just as internal boot devices have been added in the previous Solaris OS releases. Refer to the Solaris Express Developer Edition OS installation guide for details. You should consider the following points while enabling a fabric boot device.

- Minimize interference to boot devices through the following actions:
 - Ensuring the boot device is not an overly subscribed target or LUN
 - Avoiding installation of applications and software on target or LUN
 - Reducing physical distance between host and fabric device, as well as the number of hops
- Remove the boot disk from volume manager control prior to beginning the fabric boot procedure.
- Ensure that the latest HBA Fcode and drivers are loaded for the install HBAs on the system.
- If multipathing is desired on the boot device, use the `stmsboot` command as described in [Chapter 4, “Configuring Multipathing Software.”](#)

Configuring Virtual Fibre Channel Ports

This chapter provides the steps used to configure N Port ID Virtualization (NPIV) ports, also known as virtual Fibre Channel ports.

- [“What is NPIV?” on page 53](#)
- [“Limitations of NPIV” on page 53](#)
- [“Creating NPIV Ports” on page 54](#)

What is NPIV?

NPIV is a Fibre Channel facility that enables one Fibre Channel adapter to have many N Port IDs. Each N Port has a unique identity (port WWN and node WWN) on the SAN and can be used for zoning and LUN masking. Soft zoning, which you can use to group ports together by port WWN, is the preferred method of zoning.

Limitations of NPIV

Here is the list of NPIV limitations when used to virtualize Fibre Channel ports:

- NPIV ports may not be used for booting.
- NPIV ports are best used in SANs with a relatively small number of ports, either virtual or physical. Also, some targets may not have enough resources to process the large number of ports that NPIV can create. This limitation exists because processing state change notifications (SCN) on the SAN takes significant time if there are a large number of ports on the SAN. You can work around this limitation on a large SAN by using zoning, which can limit the number of visible ports.
- MPxIO can be used with NPIV, although you should ensure that different paths are physically redundant.
- NPIV is supported only in a Fabric topology. It is not supported in an FC-AL or point-to-point topology.

- Not all hardware supports NPIV. Both switches and HBAs (although not targets) must support NPIV in a SAN. By specification, HBAs should support up to 255 virtual ports, although this is defined by the resources on the switch. Switches may have to be updated to the latest firmware levels for NPIV support.

Creating NPIV Ports

The `fcadm(1M)` command is available to configure NPIV for non-virtualized environments. For xVM, see the `virsh(1M)` command documented in the Virtualization System Administration Guide.

The `fcinfo(1M)` and `fcadm(1M)` status commands are available to determine the status of NPIV ports, whether the ports are created from `fcadm(1M)` or by virtualization specific commands, such as `xm(1M)`. The commands also report the relationship between the physical port and the virtual ports hosted on that port.

Other Fibre Channel commands, such as `luxadm(1M)` and `cfgadm(1M)`, report NPIV information, although there is no distinction made between virtual and physical ports.

▼ To Create an NPIV Port

Before You Begin Each virtual port must have a port and node name. The port name must be unique on the SAN. You can assign names manually or use the built-in random WWN generator. If you attempt to register duplicate names, most switches will report an error status on the newly registered WWN, and the switch will not register the new WWN.

For more information on acceptable name formats, refer to the T11 standard: Fibre Channel Framing and Signaling (FC-FS 2).

If you try to create an NPIV port on an HBA that does not support NPIV, an error will occur. If you try to create an NPIV port on an HBA that supports NPIV, but it is attached to a switch which does not support NPIV, the port will be created with an offline status. The status will be reported in the `fcinfo(1M)` output.

1 Become superuser.

2 Create an NPIV port:

```
# fcdm create-npiv-port -p Virtual_Port_WWN -n Virtual_Node_WWN PhysicalPort_port_WWN
```

Without the `-p` and `-n` options, a random WWN will be assigned for the virtual port and virtual node, respectively.

Example 6-1 Creating an NPIV Port

The following example creates an NPIV port on a physical HBA port with a WWN of 210000e08b170f1c, a virtual port WWN set to 2000000000000001, and a virtual node WWN set to 2100000000000001.

```
# fcadm create-npiv-port -p 2000000000000001 -n 2100000000000001 210000e08b170f1c
```

▼ To Delete an NPIV Port

Before You Begin You can use the `fcinfo hba-port` command to display the current WWN values for the NPIV ports.

1 Become superuser.

2 Delete an NPIV port:

```
# fcadm delete-npiv-port -p Virtual_Port_WWN -n Virtual_Node_WWN PhysicalPort_port_WWN
```

Example 6-2 Deleting an NPIV Port

The following example deletes an NPIV port on a physical HPA port with a WWN of 210000e08b170f1c.

```
# fcadm delete-npiv-port -p 2000000000000001 -n 2100000000000001 210000e08b170f1c
```

▼ To Report Status on NPIV Ports

1 Become superuser.

2 Report status on the currently configured NPIV ports:

```
# fcinfo hba-port
```

The NPIV Port List shows the lists of currently configured NPIV ports.

Example 6-3 Reporting Status on NPIV Ports

The following example shows that HBA port 210000e08b84f7eb has one virtual port.

```
# fcinfo hba-port
HBA Port WWN: 210000e08b84f7eb
  Port Mode: Initiator
  Port ID: 10100
```

```
OS Device Name: /dev/cfg/c7
Manufacturer: QLogic Corp.
Model: 375-3294-01
Firmware Version: 04.04.00
FCode/BIOS Version: BIOS: 1.4; fcode: 1.11; EFI: 1.0;
Serial Number: 0402F00-0549112895
Driver Name: qlc
Driver Version: 20080430-0.00
Type: N-port
State: online
Supported Speeds: 1Gb 2Gb 4Gb
Current Speed: 4Gb
Node WWN: 200000e08b84f7eb
Max NPIV Ports: 63
NPIV port list:
  Virtual Port1:
    Node WWN: 1110000000000000
    Port WWN: 1210000000000000
```


Configuring FCoE Ports

This chapter provides the steps used to configure FCoE ports hosted on normal Ethernet interfaces. This chapter does not apply to hardware FCoE ports on Converged Network Adapters (CNA).

The following topics are discussed:

- “What is FCoE?” on page 57
- “Limitations of FCoE” on page 57
- “Configuring FCoE Ports” on page 58

What is FCoE?

FCoE is a new T11 standard that transports encapsulated Fibre Channel frames over Enhanced Ethernet. FCoE is designed to enable network convergence and cost-effective SAN expansion in data centers.

Limitations of FCoE

The Solaris FCoE initiator is a pure software implementation that is designed to work with normal Ethernet controllers. However, you must comply with the following limitations of Solaris FCoE implementations:

- FCoE ports cannot be used for booting.
- FCoE ports cannot be configured in xVM or LDom guest operating systems.
- FCoE is supported in fabric and point-to-point topologies.

FCoE is not supported on all hardware. FCoE works with Ethernet controllers that support 802.3x PAUSE and jumbo frames, and have a GLDv3 driver.

Configuring FCoE Ports

The `fcadm(1M)` command is available to configure FCoE ports. The `fcinfo(1M)` and `fcadm(1M)` status commands are available to determine the status of FCoE ports. The commands also report the relationship between the Ethernet interface and the FCoE port hosted on that interface.

Other Fibre Channel commands, such as `luxadm(1M)` and `cfgadm(1M)`, report FCoE information, although there is no distinction made between FCoE and native FC ports.

▼ To Create an FCoE Port

Before You Begin Before you begin this procedure, you must perform the following tasks:

- Enable the 802.3x (also called PAUSE) setting on the Ethernet interface. This setting ensures a lossless Ethernet transport.
- Enable jumbo frames (greater than 2.5 Kbytes) on the Ethernet interface. A Fibre Channel data frame can be as large as 2136 bytes.

These settings can vary for different Ethernet hardware and drivers. In most cases, you must modify the `driver.conf` file of the Ethernet interface and then reboot. See the `driver.conf` file for your Ethernet interface for details on how to enable these features.

Each virtual port must have a port and node name. The port name must be unique on the SAN. You can assign names manually or use the built-in WWN generator. If you attempt to register duplicate names, the switch will report an error status on the newly registered WWN, and the switch will not register the new WWN. For more information on acceptable name formats, refer to the T11 standard: Fibre Channel Framing and Signaling (FC-FS 2).

If you try to create an FCoE port on a network interface that does not support FCoE (for example, `xnf` in `xVM` or `VNET` in `LDom`), an error occurs and the FCoE port is not created.

1 Become superuser.

2 Create an FCoE port.

```
# fcdm create-fcoe-port -p Port_WWN -n Node_WWN Ethernet_Interface
```

If the selected Ethernet interface does not support Multiple Unicast Address, you are prompted to explicitly enable promiscuous mode on that interface.

```
# fcdm create-fcoe-port -f Ethernet_Interface
```

Example 7-1 Creating an FCoE Port

The following example creates an FCoE port on Ethernet interface `nxge0` with port WWN and node WWN generated automatically.

```
# fcadm create-fcoe-port nxge0
```

▼ To Delete an FCoE Port

Before You Begin You can use the `fcadm list-fcoe-ports` command to display the Ethernet interfaces hosting the FCoE ports.

1 Become superuser.

2 Delete an FCoE port.

```
# fcadm delete-fcoe-port network_interface
```

Example 7-2 Deleting an FCoE Port

The following example deletes an FCoE port on Ethernet interface `nxge0`.

```
# fcadm delete-fcoe-port nxge0
```

▼ To Report Status on FCoE Ports

1 Become superuser.

2 Report the status on the currently configured FCoE ports.

```
# fcinfo hba-port -e
```

This command retrieves a list of FC specific information for all FCoE ports in the system.

```
# fcadm list-fcoe-ports
```

This command retrieves a list of FCoE specific information for all FCoE ports in the system.

Example 7-3 Reporting Status on FCoE Ports

The following example reports FC and FCoE specific information for all FCoE ports.

```
# fcinfo hba-port -e
HBA Port WWN: 200000144fc1f5c8
      Port Mode: Initiator
```

```
Port ID: 9a0042
OS Device Name: /dev/cfg/c6
Manufacturer: Sun Microsystems, Inc.
Model: FCoE Virtual FC HBA
Firmware Version: N/A
FCode/BIOS Version: N/A
Serial Number: N/A
Driver Name: SunFC FCoEI v20090422-1.00
Driver Version: v20090422-1.00
Type: N-port
State: online
Supported Speeds: 1Gb 10Gb
Current Speed: 10 Gb
Node WWN: 100000144fc1f5c8
```

```
# fcdm list-fcoe-ports
```

```
HBA Port WWN: 200000144fc1f5c8
Port Type: Initiator
MAC Name: nxge0
MTU Size: 9194
Primary MAC Address: 00144fc1f5c8
Current MAC Address: 0efc009a0042
Promiscuous Mode: Off
```

Configuring SAS Domains

This chapter provides overviews of SAS Domain considerations, SAS device discovery, and SAS boot device configuration.

SAS Multipathing Considerations

- SAS multipathing is supported with the Solaris Express Developer Edition OS release when using the bundled `mpt (7d)` driver.
- SAS expanders are not supported with the Solaris Express Developer Edition OS release.
- LUN masking is not supported with SAS-attached devices in the Solaris Express Developer Edition OS release.
- Turn off power management on hosts connected to the SAS Domain in order to prevent unexpected results as one server attempts to power down a device while another attempts to gain access. See `pmconfig (1m)` and `power.conf (4)` for details about power management.

Dynamic Discovery of SAS Devices

Adding and removing SAS devices is performed dynamically when using the `mpt (7d)` driver. You are no longer required to edit the `sd.conf` file to enable your system to detect specific targets and LUNs attached with the `mpt` driver. See `mpt (7d)` for more information.

You are no longer required to edit the `sd.conf` file to enable your system to detect specific targets and LUNs attached with the `mpt` driver.

If you add or remove a device in your SAS domain, messages will be written to the `/var/adm/messages` file indicating presence or removal. If a device has been added, it is visible to and available by the `format (1m)` command.

Configuring SAS Boot Devices

The Solaris Multipathed IO software allows Solaris Express Developer Edition OS hosts to boot from a multipathed SAS device, or from a SATA device connected to a SAS controller. Only the bundled `mpt` driver is supported for SAS multipathing in the Solaris Express Developer Edition OS release.

Some SAS and SATA devices might not support multipathing but will still function as non-multipathed devices. See <http://sunsolve.sun.com> for more information.

Configuring IPFC SAN Devices

Configuring Internet Protocol over Fibre Channel (IPFC) on the host system describes host recognition of IPFC devices and implementation of IP over FC in a SAN. The IPFC driver is based on RFC 2625 and allows IP traffic to run over FC.

This chapter contains the following topics:

- “IPFC Considerations” on page 63
- “Invoking and Configuring IPFC” on page 66

IPFC Considerations

The following table shows the supported features available for IPFC.

TABLE 9-1 IPFC (NFS/NAS and SNDR)

Feature	Supported
Cascading	Yes, with fabric zones only
Zone type	Fabric zone with the HBA configured as an F-port point-to-point connection
Maximum number of device ports per zone	253

The following restrictions apply:

- IPFC is not supported on 1 Gbit Sun switches.
- Promiscuous mode is not supported. The snoop(1M) utility cannot be used.
- Multicasting is supported through broadcasting only.
- Network cards using IPFC cannot be used as routers. The /etc/not_router file must be present on the host.

- Any standard network commands can be used after IPFC is attached. These commands (telnet, ping, or ftp) are used in this environment in the same way as in an Ethernet setup.

Determining Fibre Channel Adapter Port Instances

This section explains how to configure the desired host system for IPFC. It includes the procedures to determine the port instance and to plumb an IPFC instance.

▼ To Determine Port Instances

1 Determine the HBA PCI adapter slot and the I/O board PCI slot.

You need this information to perform the calculation in “[Determining Fibre Channel Adapter Port Instances](#)” on page 64.

For example, assume you have an array with an HBA card located in PCI adapter slot 5, and the PCI adapter is in slot 1 of the I/O board.

2 Determine the instance number.

a. Search for the fp driver binding name in the /etc/path_to_inst file.

Note – Determine the correct entry by finding the hardware path described in your server hardware manual or Sun System Handbook. The Sun System Handbook is available at http://sunsolve.sun.com/handbook_pub/.

b. Narrow the search by using the I/O board and slot information from Step 1.

Note – The following method of deriving the Solaris Express Developer Edition OS device path of an HBA from its physical location in server might not work for all Sun server hardware.

i. Multiply the PCI adapter slot number by the number of adapter ports.

For example, if the HBA has two ports, multiply by 2. Using the array with an HBA in the PCI adapter slot 5, multiply 5 by 2 to get 10.

ii. Add the PCI adapter I/O board slot number to the number derived in Step i.

Using an HBA in PCI adapter slot 5 and PCI slot 1 of the I/O board, add 1 to 10 for a sum of 11.

iii. Convert the number derived in Step ii to hexadecimal.

The number 11 converts to “b” in hexadecimal.

iv. Search for the fp entry with `pci@hex` where `hex` is the number you derived in Step iii.

The following table shows the elements of the device path for a PCI single FC network adapter device that has the following path:

```
"/pci@b,2000/SUNW,qlc@2/fp@0,0" 7 "fp"
```

EntryItem	EntryValue
Physical Name	/pci@b,2000/SUNW,qlc@2/fp@0,0
Instance Number	7
Driver Binding Name	fp

c. Manually plumb each FP instance.

Use the `ifconfig interface-number plumb` command. In this example, the value of `interface-number` is `fcip7`.

```
# ifconfig fcip7 plumb
```

When the command is successful, a message appears on both the console and in the messages file. For example:

```
Sep 13 15:52:30 bytownite ip: ip: joining multicasts failed (7) on fcip0 -
will use link layer brocasts for multicast
```

▼ To Plumb an IPFC Instance

Each FP instance on the system has an entry in `/dev/fc`. If HBAs have been removed, some stale links might exist. Use this procedure to load and plumb IPFC.

1 For each entry in `/dev/fc`, issue a `luxadm -e dump_map` command to view all the devices that are visible through that HBA port:

```
# luxadm -e dump_map /dev/fc/fp0
Pos  Port_ID Hard_Addr Port WWN          Node WWN          Type
0    610100 0          210000e08b049f53 200000e08b049f53 0x1f (Unknown Type)
1    620d02 0          210000e08b02c32a 200000e08b02c32a 0x1f (Unknown Type)
2    620f00 0          210000e08b03eb4b 200000e08b03eb4b 0x1f (Unknown Type)
3    620e00 0          210100e08b220713 200100e08b220713 0x1f (Unknown Type,Host Bus Adapter)
# luxadm -e dump_map /dev/fc/fp1
No FC devices found. - /dev/fc/fp1
```

2 Based on the list of devices, determine which destination HBAs are visible to the remote host with which you want to establish IPFC communications.

In the example for this procedure, the destination HBAs have port IDs `610100` and `620d02`. The originating HBA's port ID is `620e00`.

- 3 List the physical path of the originating HBA port from which you can see the destination HBA port, where *originating-hba-link* is a variable for the link determined in Step 2.

```
# ls -l /dev/fc/fp originating-hba-link
```

In the following example, 0 is the number for the *originating-hba-link*:

```
# ls -l /dev/fc/fp 0
lrwxrwxrwx 1 root root 51 Sep 4 08:23 /dev/fc/fp0 ->
../devices/pci@8,600000/SUNW,qlc@1/fp@0,0:devctl
```

- 4 Search the physical path identified in Step 3.

You must remove the leading `../devices` from the path name output. For example

```
# grep pci@8,600000/SUNW,qlc@1/fp@0,0 /etc/path_to_inst
"/pci@8,600000/SUNW,qlc@1/fp@0,0" 0 "fp"
```

- 5 Determine the `fp` instance for the originating HBA port from the output of the command in Step 4.

The instance number precedes “`fp`” in the output. In the following example output, the instance number is 0.

```
"/pci@8,600000/SUNW,qlc@1/fp@0,0" 0 "fp"
```

- 6 Use the instance number from Step 5 to load IPFC and plumb the IPFC interface.

In this example, the instance is 0.

```
# ifconfig fcip 0 plumb
```

Invoking and Configuring IPFC

Immediately upon installation, start IPFC manually with the `ifconfig` command. You can configure the host so that on subsequent reboot, the IPFC network interface starts automatically. This section describes the procedures to start a network interface manually and to configure the host for automatic plumbing upon reboot.

▼ To Start a Network Interface Manually

Use this procedure when you want to plumb IPFC with specific netmask values and get the IPFC interface up and running.

1 Use the `ifconfig` command with the appropriate interface.

Ask your network administrator for an appropriate IP address and netmask information. For example, to enable an IPFC interface associated with `fp` instance `0` and an IP address of `192.9.201.10`, type:

```
# touch /etc/notrouter
# ifconfig fcip0 inet 192.9.201.10 netmask 255.255.255.0 up
```

The `ifconfig` command is described in more detail in the `ifconfig(1M)` man page.

2 Use the command `ifconfig -a` to verify that the network is functioning.

The output of `ifconfig -a` should look like this:

```
lo0: flags=1000849<UP,LOOPBACK,RUNNING,MULTICAST,IPv4> mtu 8232 index 1
inet 127.0.0.1 netmask ffffffff
fcip0: flags=1001843<UP,BROADCAST,RUNNING,MULTICAST,MULTI_BCAST,IPv4> mtu 1500 index 2
inet 192.9.201.10 netmask ffffffff broadcast 192.9.201.255 ether 0:e0:8b:1:3c:f7
hme0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu 1500 index 3
inet 192.9.200.70 netmask ffffffff broadcast 192.9.200.255
ether 8:0:20:fc:e9:49
```

▼ To Configure the Host for Automatic Plumbing Upon Reboot

Each network interface must have an `/etc/hostname.interface` file defining the name of the IP address associated with it. For example, IPFC network interface `fcip0` has a file named `/etc/hostname.fcip0`.

1 Manually create a `/etc/hostname.interface` file that contains a single line that identifies the host name or interface IP address.

2 Make any additional entries to the `/etc/inet/hosts` file.

The Solaris Express Developer Edition OS installation program creates the `/etc/inet/hosts` file with minimum entries. You must manually make additional entries with a text editor. See the `hosts(4)` man page for additional information.

The `/etc/inet/hosts` file contains the hosts database. This file contains the host names and the primary network interface IP addresses, as well as the IP addresses of other network interfaces attached to the system and of any other network interfaces that the machine must know about.

The following example shows an `etc/inet/host` file.

```
127.0.0.1    localhost    loghost
192.9.200.70 sun1        #This is the local host name
192.9.201.10 fcip0      #Interface to network 192.9.201.10
```

- 3 Edit the `/etc/nsswitch.conf` file so that all uncommented entries have the word `files` before any other name service.**

The `/etc/nsswitch.conf` specifies which name service to use for a particular machine. The following code shows an example of an `/etc/nsswitch.conf` file.

```
hosts: files nis
```

Sun Solaris Express Developer's Edition OS x86/x64 Fibre Channel Operating System Booting Procedure

In the Solaris Express Developer's Edition, x86/x64 Operating System (OS), Fibre Channel (FC) devices now appear in the disk selection portion of Solaris SunInstall. An interactive Solaris Express Developer's Edition installation is now required due to the potentially large number of devices that might appear in the disk selection portion of Solaris SunInstall.

This chapter describes how to manually install the Solaris Express Developer's Edition. This OS includes Sun StorageTek SAN Foundation software with the Sun 2-Gbit and 4-Gbit Fibre Channel (FC) HBA drivers, and the Sun StorageTek Traffic Manager software.

Note –

Sun 1 Gb HBAs do *not* support booting over a SAN on x86 machines. Sun 4 Gb HBAs do support booting over a SAN on x86 machines. Most Sun 2 Gb HBAs do support booting over a SAN on x86 machines, except for the following:

- Sun StorageTek 2 Gb Enterprise Class Dual-Port Fibre HBA, SG-XPCI2FC-QF2-Z
- Sun StorageTek 2 Gb PCI Dual-Port Fibre HBA, SG-XPCI2FC-QF2
- Sun StorageTek 2 Gb FC PCI Single-Channel Network Adapter, X6767A
- Sun StorageTek 2 Gb FC PCI Dual-Channel Network Adapter, X6768A
- Sun StorageTek 2 Gb FC PCI-X Enterprise Single-Port HBA, SG-XPCI1FC-QL2
- Sun StorageTek 2 Gb FC PCI Single-Port Host Adapter, SG-XPCI1FC-QF2

The Sun Solaris Express Developer's Edition x64/x86 OS can be installed from a CD-ROM, DVD-ROM, or network.

This chapter contains the following topics:

- [“Setup Requirements” on page 70](#)
- [“Installation Overview” on page 70](#)
- [“Installation Procedure” on page 71](#)
- [“Post-Installation Patches for the Sun Solaris Express Developer's Edition OS” on page 78](#)

Setup Requirements

You must have the following items for installation.

- Solaris Express Developer's Edition
- 10/100/1000 Mbit/sec Ethernet network for network-based OS installation
- One of the following for installation support
 - Jump-start server for a network-based OS installation
 - CD-ROM or DVD-ROM drive for a CD-ROM or DVD-ROM based OS installation

Installation Overview

In order to successfully complete the installation, use an FC-based device during the disk selection portion of the installation. At the end of the interactive installation, you must select Manual Reboot.

After OS installation and before rebooting, gather the configuration information by issuing the `luxadm` command on the newly installed logical unit number (LUN). The `luxadm` command output provides a map from `c#t#d#` to the array WWN and LUN number. Record the HBA World Wide Name (WWN) and array WWN port information.

During rebooting, use the WWN and LUN number to set the HBA BIOS for each HBA to be used to boot from each LUN on the same array. Change the system BIOS to perform booting from CD-ROM or Network to disk.

Note – Console access is required for HBA and system BIOS changes.

In addition, note the following during the Sun Solaris OS installation.

- If using a custom disk layout, do not remove the overlap (`s2`) partition. The x86 boot installer has a dependency on this partition.
- By default, the Sun StorageTek Traffic Manager software manages the FC boot devices with multiple paths to a single boot device.

Installation Procedure

▼ To Summarize the Procedure

- 1 **Install the HBA hardware.**
Follow the instructions in the appropriate Sun HBA installation guide found at <http://docs.sun.com/app/docs?q=StorageTek+Host+Bus+Adapter>.
- 2 **Install the Sun Solaris OS as described in the next section, “Installation Procedure” on page 71**

▼ To Perform a CD-ROM, DVD ROM, or Network-Based OS Installation

After installing the HBA , perform the following steps for an interactive network-based installation of the Sun Solaris Express Developer's Edition x86/x64 OS.

For more information, refer to *Solaris Express Installation Guide: Basic Installations*, 819-2393 at <http://docs.sun.com/>.

- 1 **If you are installing from a CD-ROM or DVD-ROM rather than over a network, insert the Sun Solaris Software into the CD-ROM or DVD-ROM drive.**
- 2 **Upon initial power up, provide the system BIOS and configure it to boot from either the network or CD-ROM (DVD-ROM) as applicable.**

See [Figure 10–1](#) for an example of one of several BIOS setup screens.

Note – The system BIOS setup screens are different for each system BIOS vendor.

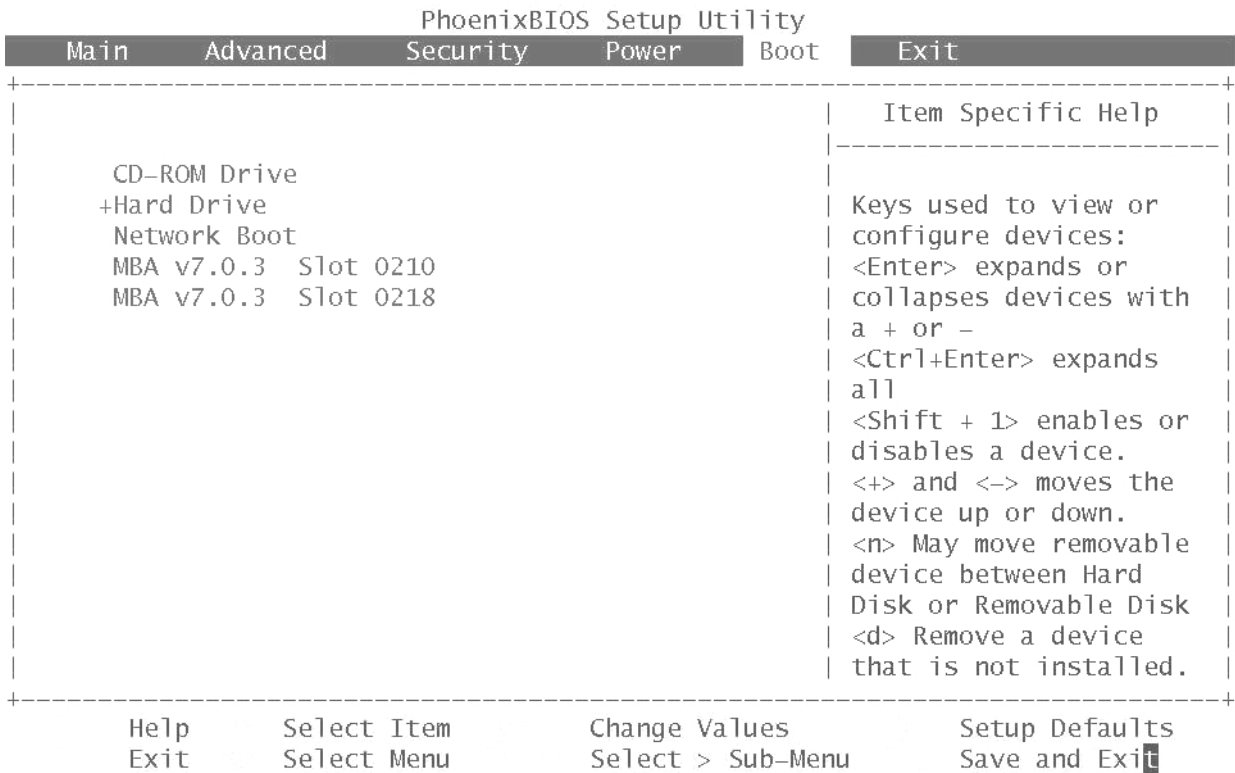


FIGURE 10-1 System BIOS for Selecting Disk, CD-ROM, or Network Boot Device Priority

- 3 Install the Sun Solaris OS using the Solaris Interactive Installation (SunInstall). See [Figure 10-2](#) for an example of one of several SunInstall screens.


```

Select Disks

On this screen you must select the disks for installing Solaris software.
Start by looking at the Suggested Minimum field; this value is the
approximate space needed to install the software you've selected. Keep
selecting disks until the Total Selected value exceeds the Suggested Minimum
value.
NOTE: ** denotes current boot disk

Disk Device                                     Available Space
=====
^ [ ] c0t500060E800000000000000EBB100000189d0      1020 MB
x [ ] c0t500060E800000000000000EBB10000018Ad0      1020 MB
x [ ] c0t500060E800000000000000EBB10000018Bd0      1020 MB
x [ ] c0t500060E800000000000000EBB10000018Cd0      1020 MB
x [x] c0t600015D0002028000000000000001142d0      10208 MB (F4 to edit)
x [ ] c0t600015D0002028000000000000001146d0      10208 MB
x [ ] c0t600015D000202800000000000000114Ad0      10208 MB
v [ ] c0t600015D000202800000000000000114Ed0      10208 MB

Total Selected: 10208 MB
Suggested Minimum: 3574 MB

F2 Continue  F3 Go Back  F4 Edit  F5 Exit  F6 Help

```

FIGURE 10-2 Selecting an FC-based Boot Disk During Disk Selection

- 4 During installation, select a desired array and its associated LUN.
- 5 Select Manual Reboot at the end of the SunInstall utility screens and then verify your selections to start the Sun Solaris OS installation.
- 6 Before rebooting after the installation completes, issue the `luxadm display` command on the boot LUN.

See Figure 10-3.

```

# luxadm display /dev/rdisk/c0t600015D0002028000000000000001142d0s2
DEVICE PROPERTIES for disk: /dev/rdisk/c0t600015D0002028000000000000001142d0s2
Vendor:      SUN
Product ID:  SE6920
Revision:    0202
Serial Num:  00500057
Unformatted capacity: 10240.000 MBytes
Read Cache:  Enabled
  Minimum prefetch: 0x0
  Maximum prefetch: 0xffff
Device Type: Disk device
Path(s):

/dev/rdisk/c0t600015D0002028000000000000001142d0s2
/devices/scsi_vhci/disk@g600015d0002028000000000000001142:c,raw
Controller   /dev/cfg/c4
Device Address      213600015d207200,0

```

```

Host controller port WWN          210100e08b206812
Class                            primary
State                            ONLINE
Controller                        /dev/cfg/c11
Device Address                    213600015d207200,0
Host controller port WWN          210100e08b30a2f2
Class                            primary
State                            ONLINE

```

```

QLogic FastUTIL
+-----Selected Adapter-----+
Adapter Type I/O Address Slot Bus Device Function
QLA2462      3400          02  03  01    1
+-----+
+-----Adapter Settings-----+
BIOS Address:                CF000
BIOS Revision:                1.05
Adapter Serial Number:       A04712
Interrupt Level:              5
Adapter Port Name:            210100E08B206812
Host Adapter BIOS:            Enabled
Frame Size:                   2048
Loop Reset Delay:              5
Adapter Hard Loop ID:         Disabled
Hard Loop ID:                  0
Spinup Delay:                  Disabled
Connection Options:           2
Fibre Channel Tape Support: Enabled
Data Rate:                     2
+-----+
Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

```

FIGURE 10-3 luxadm display Command and Output

The following luxadm command output from the example in the figure can be used to map the MPxIO based c#t#d# to the HBA WWN and array WWN:

- MPxIO c#t#d# = c0t600015d000202800000000000000001142d0
- Array WWN = 213600015d207200, LUN 0
- HBA WWNs = 210100e08b206812 and 210100e08b30a2f2

- 7 **During the reboot process, watch the monitor so you can enter the HBA #1 BIOS screen and specify the boot device to be the FC LUN onto which you just installed the Sun Solaris OS.**

Follow this step for each HBA to be used for multipathing, and specify the boot device to be the FC LUN on which you installed the Sun Solaris OS. See [Figure 10–4](#) and [Figure 10–5](#).

 - **For the QLogic HBA BIOS, perform the following steps.**
 - a. **During host rebooting, press Control-Q to display the HBA BIOS screen.**
 - b. **Select the HBA that you want to enable as the boot device and enable boot.**
 - c. **Configure the boot device:**
 - i. **Select Configuration Settings.**
 - ii. **Select Selectable Boot Settings.**
 - iii. **Make sure that Selectable Boot is set to enable.**

In this menu, you can select the boot device/LUN by the array WWPN.
 - iv. **Save and then exit the HBA BIOS screen.**
 - **For the Emulex HBA BIOS, perform the following steps.**
 - a. **During host rebooting, press Alt-E to display the HBA BIOS screen.**
 - b. **Select the HBA that you want to enable as the boot device and enable boot.**
 - c. **Select Configure Boot Devices.**
 - d. **Select a boot entry.**
 - e. **Select the WWPN of the desired boot device.**
 - f. **Type the LUN number.**
 - g. **Select the boot LUN.**
 - h. **Select Boot Device with the Array WWPN.**
 - i. **Save and exit the HBA BIOS screen.**

```
QLogic Fast!UTIL

-----Selected Adapter-----
Adapter Type I/O Address Slot Bus Device Function
QLA2462      3400          02  03  01    1

-----Adapter Settings-----

BIOS Address:           CF000
BIOS Revision:          1.05
Adapter Serial Number:  A04712
Interrupt Level:        5
Adapter Port Name:      210100E08B206812
Host Adapter BIOS:      Enabled
Frame Size:             2048
Loop Reset Delay:       5
Adapter Hard Loop ID:   Disabled
Hard Loop ID:           0
Spinup Delay:           Disabled
Connection Options:     2
Fibre Channel Tape Support: Enabled
Data Rate:              2

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit
```

FIGURE 10-4 HBA BIOS Screen for an HBA WWN

```

QLogic FastUTIL
=====Selected Adapter=====
| Adapter Type I/O Address Slot Bus Device Function |
| QLA2462      3000      02  03  01  0 |
=====

=====Selectable Boot Settings=====
| Selectable Boot: Enabled |
| (Primary) Boot Port Name,Lun: 213600015D207200, 0 |
| Boot Port Name,Lun: 213600015D207200, 1 |
| Boot Port Name,Lun: 213600015D207200, 2 |
| Boot Port Name,Lun: 213600015D207200, 3 |
| Press "C" to clear a Boot Port Name entry |
=====

Use <arrow keys> and <Enter> to change settings, <Esc> to exit

```

FIGURE 10-5 HBA BIOS Screen for Enabling Boot From the FC LUN

Note –

The figure shows the following modifications:

- Selectable boot = Enabled
- ARRAY WWN = 213600015d207200
- ARRAY LUN = 0
- HBA WWN = 210100e08b206812

- 8 Repeat the appropriate modifications for all HBAs and all LUNs.
- 9 Type the system BIOS per the vendor's access method and specify the boot device to be the FC LUN on which you installed the Solaris Express Developer's Edition OS. The following figure shows an example of one screen in a vendor's system BIOS.

```

PhoenixBIOS Setup Utility
-----
Main      Advanced  Security  Power    Boot     Exit
-----
- Hard Drive
  HA 1 PortID 180F00 SE6920
  HA 2 PortID 180F00 SE6920
  #220 ID00 LUN0 SEAGATE ST373307LC
  Bootable Add-in Cards
  CD-ROM Drive
  Network Boot
  MBA v7.0.3 Slot 0210
  MBA v7.0.3 Slot 0218
- Removable Devices
  (F1) Help (F9) Setup Defaults
  (Esc) Exit (F10) Save and Exit

Item Specific Help
-----
Keys used to view or
configure devices:
<Enter> expands or
collapses devices with
a + or -
<Ctrl+Enter> expands
all
<Shift + I> enables or
disables a device.
<+> and <-> moves the
device up or down.
<n> May move removable
device between Hard
Disk or Removable Disk
<d> Remove a device
that is not installed.
-----
F1 Help ^v Select Item -/+ Change Values F9 Setup Defaults
Esc Exit <-> Select Menu Enter Select > Sub-Menu F10 Save and Exit

```

FIGURE 10-6 Example System BIOS for Enabling Booting from an FC LUN

- 10 Reboot to the newly installed Solaris Express Developer's Edition OS using the FC LUN specified in the system BIOS.

Post-Installation Patches for the Sun Solaris Express Developer's Edition OS

The following patches are recommended for the OS:

- 120223-xx Emulex-Sun LightPulse FC adapter driver
- 119131-xx FC device drivers
- 119716-xx MPxIO
- 119471-xx luxadm

These patches can be found on the SunSolveSM web site at <http://sunsolve.sun.com>.

Persistent Binding for Tape Devices

To simplify management of servers in SAN-based data-centers, the Sun StorageTek™ SAN Foundation software stack in the Solaris OS dynamically detects devices in a SAN and builds associated `/dev` tree entries without requiring you to edit configuration files.

In most cases, this process greatly simplifies SAN management. However, for tape devices, you might like the ability to explicitly specify how that `/dev` entries are created and to ensure the `/dev` entries are identical across multiple servers on a SAN. This chapter describes how you can specify this tape binding in the Solaris OS while retaining the benefits of automatic discovery for disk-based devices.

Persistent Binding Overview

The `/dev/rmt` directory contains links to physical devices under `/devices` for tape devices. Each tape LUN seen by the system is represented by 24 minor nodes in the form of `/dev/rmt/N`, `/dev/rmt/Nb`, and `/dev/rmt/Nbn`, where N is an integer counter starting from 0. This number is picked by `devfsadm` during enumeration of new devices. Every new tape logical unit number (LUN) found by `devfsadm` gets the next available number in `/dev/rmt`.

Because the `/dev/rmt` name depends on the order in which devices appear in the device tree, it changes from host to host. For a given tape drive that is seen by two or more different hosts, the `/dev/rmt` link can be different on each of these hosts. This difference can cause problems for the most common usage of Symantec (VERITAS) NetBackup (SSO option). Also, if the drive is replaced, the links change unless the vendor provides a way to retain the port World-Wide Name (*PWWN*) of the drive.

Creating Tape Links

The `/etc/devlink.tab` file is called the default device table file. It specifies rules that `devfsadm` uses to create links in the `/dev` directory. Type `man devlinks` on the command-line interface (CLI) for more information. This file does not contain any entries for tapes because `devfsadm` is already able to create links for tape drives, but rules can be added that will modify the default behavior for creating tape links.

For any tape drive visible to the Sun Solaris OS but not specified in the `devlink` file, `devfsadm` automatically assigns a minor node number starting from 0. These minor node numbers will conflict with any lower numbers assigned manually in `/etc/devlink.tab`, so be sure to assign numbers there that are high enough to avoid conflicts.

This approach can easily lead to duplicate links in `/dev/rmt`. Any tapes discovered before entries were specified in `/etc/devlink.tab` have automatically created links. When entries are added and `devfsadm` is run, the original links remain in `/dev/rmt`, resulting in duplicate links. To remove the original links in `/dev/rmt`, run the `rm /dev/rmt/*` command before running `devfsadm`.

This approach cannot be used with multiple-port tape drives that are attached to multiple HBA ports. If multiple HBA ports are attached to the same tape LUN, the system detects two tape drives instead of one. The one that appears last in the `prtconf` output gets the link generated by the `/etc/devlink.tab`.

The following example shows a sample entry for tape in the `devlink.tab` file.

```
type=ddi_byte:tape;addr=PWWN,LUN-number; rmt/rmt-number\M0
```

Change the `rmt #` to whatever `/dev/rmt/N` is required. Then change the PWWN and LUN to match the desired tape device. You can obtain this value by running the `ls -l` command on the existing `/dev/rmt/` link as shown below.

```
# ls -l /dev/rmt/4
lrwxrwxrwx 1 root root 69 Oct 6 14:57 /dev/rmt/4 ->
../../../../devices/pci@1f,700000/SUNW,qlc@2/fp@0,0/st@w5005076300617717,0:
```

If, for example, you wanted the `/dev/rmt/` number to be 40, you would create an entry in `/etc/devlink.tab` like the following example:

```
# type=ddi_byte:tape;addr=w5005076300617717,0; rmt/40\M0
```

You can then add this line to the `devlink` file on every Solaris server on the SAN that uses this drive so that it always appears as minor node 40.

▼ Configuration Steps

- 1 **Create the entries in `/etc/devlink.tab` as described in “Creating Tape Links” on page 80.**

If `devfsadm` has previously discovered the devices, you must determine the device address by running the `ls -l` command on the existing link.

Note – Be sure to assign `/dev/rmt/ N` numbers to avoid conflicts with any automatically configured devices, as described above.

- 2 **Remove existing links from `/dev/rmt` by running the `rm /dev/rmt/*` command.**
- 3 **Run `devfsadm`.**

This command creates new links as per the entries in `/etc/devlink.tab` in addition to automatically creating links for any unspecified devices.

Manual Configuration for Fabric-Connected Devices

This appendix explains about configuring and unconfiguring the fabric devices in the Solaris Express Developer's Edition OS. It explains how the visible fabric devices on a host are detected and configured with and without enabling the Sun StorageTek Traffic Manager software.

This appendix contains the following sections:

- “Manual Configuration” on page 83
- “Fabric Device Node Configuration” on page 84
- “Configuring Device Nodes Without Multipathing Enabled” on page 86
- “Configuring Device Nodes With Solaris Multipathing Software Enabled” on page 89
- “Unconfiguring Fabric Devices” on page 93

Manual Configuration

In the Solaris Express Developer's Edition OS, fabric-connected devices are made available to the Solaris host. This differs from previous versions of the Sun StorageTek Traffic Manager software running on the Solaris 8 or the Solaris 9 OS. In those versions, manual configuration steps were required to achieve fabric-connected device availability.

If you want to manually configure the fabric-connected devices, use the following steps to change the default behavior on the Solaris Express Developer's Edition OS.

Note – Changing the default behavior makes all of your fabric-connected devices unavailable, which can cause problems for fabric-connected devices that are required to be available at boot time.

▼ To Configure Manually

- 1 To enable manual configuration, make sure that the following line in the `/kernel/drv/fp.conf` file does not start with an initial `#` character:

```
manual_configuration_only=1;
```

Refer to the `fp(7D)` man page and the `cfgadm_fp(1M)` man page for further information on this setting.

- 2 Reboot the host.

- 3 For each fabric-connected device that will need to be made available, follow the steps in [“Configuring Device Nodes Without Multipathing Enabled” on page 86](#) or [“Configuring Device Nodes With Solaris Multipathing Software Enabled” on page 89](#), depending on whether you are using the Sun StorageTek Traffic Manager software for multipathing.

If the original default behavior for fabric-connected devices is desired, use the following steps:

- 4 To disable manual configuration, make sure that the following line in the `/kernel/drv/fp.conf` file starts with an initial `#` character:

```
# manual_configuration_only=1;
```

Refer to the `fp(7D)` man page and the `cfgadm_fp(1M)` man page for further information on this setting.

- 5 Reboot the host.

Fabric Device Node Configuration

After you configure the hardware in your direct-attach system or SAN, you must ensure that the hosts recognize the devices. This section explains host recognition of fabric devices, also known as 24-bit FC addressing devices on the SAN. After configuring the devices, ports, and zones in your SAN, make sure that the host is aware of the devices. You can have up to 16 million fabric devices connected together on a SAN with FC support.

This section is limited to the operations required from the perspective of the Solaris Express Developer's Edition OS. It does *not* cover other aspects, such as device availability and device-specific management. If devices are managed by other software, such as a volume manager, refer to the volume manager product documentation for additional instructions.

Ensuring That LUN Level Information Is Visible

▼ To Ensure LUN Level Information is Visible

1 Use the `cfgadm` command to identify LUN level information.

If you issue the `cfgadm -al -o show_SCSI_LUN controller-ID` command immediately after a system boots up, the output might not show the Fibre Channel Protocol (FCP) SCSI LUN level information. The information does not appear because the storage device drivers, such as the `ssd` and `st` driver, are not loaded on the running system.

2 Use the `modinfo` command to check whether the drivers are loaded. After the drivers are loaded, the LUN level information is visible in the `cfgadm` output.

▼ To Detect Fabric Devices Visible on a Host

This section provides an example of the procedure for detecting fabric devices using FC host ports `c0` and `c1`. This procedure also shows the device configuration information that is displayed with the `cfgadm(1M)` command.

Note – In the following examples, only failover path attachment point IDs (`Ap_Ids`) are listed. The `Ap_Ids` displayed on your system depend on your system configuration.

1 Log in as root (`su - root`).

2 Display the information about the attachment points on the system.

```
# cfgadm -l
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	unconfigured	unknown
c1	fc-private	connected	configured	unknown

In this example, `c0` represents a fabric-connected host port, and `c1` represents a private, loop-connected host port. Use the `cfgadm(1M)` command to manage the device configuration on fabric-connected host ports.

By default, the device configuration on private, loop-connected host ports is managed by a host using the Solaris Express Developer's Edition OS.

3 Display information about the host ports and their attached devices.

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	unconfigured	unknown
c0::50020f2300006077	disk	connected	unconfigured	unknown
c0::50020f23000063a9	disk	connected	unconfigured	unknown
c0::50020f2300005f24	disk	connected	unconfigured	unknown

```

c0::50020f2300006107 disk      connected  unconfigured unknown
c1                fc-private connected  configured  unknown
c1::220203708b69c32b disk    connected  configured  unknown
c1::220203708ba7d832 disk    connected  configured  unknown
c1::220203708b8d45f2 disk    connected  configured  unknown
c1::220203708b9b20b2 disk    connected  configured  unknown

```

Note – The `cfgadm -l` command displays information about FC host ports. You can also use the `cfgadm -al` command to display information about FC devices. The lines that include a port world wide name (WWN) in the `Ap_Id` field associated with `c0` represent a fabric device. Use the `cfgadm configure` and `unconfigure` commands to manage those devices and make them available to hosts using the Solaris Express Developer's Edition OS. The `Ap_Id` devices with port WWNs under `c1` represent private-loop devices that are configured through the `c1` host port.

Configuring Device Nodes Without Multipathing Enabled

This section describes fabric device configuration tasks on a host that does not have multipathing enabled.

The procedures in this section show how to detect fabric devices that are visible on a host and to configure and make them available to a host using the Solaris Express Developer's Edition OS. The procedures in this section use specific devices as examples to illustrate how to use the `cfgadm(1M)` command to detect and configure fabric devices.

The device information that you supply and that is displayed by the `cfgadm(1M)` command depends on your system configuration.

▼ To Configure an Individual Device Without Multipathing

This sample procedure describes how to configure a fabric device that is attached to the fabric-connected host port `c0`.

- 1 Log in as root (`su - root`).
- 2 Identify the device to be configured.

```

# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected  unconfigured unknown
c0::50020f2300006077 disk      connected  unconfigured unknown
c0::50020f23000063a9 disk      connected  unconfigured unknown
c0::50020f2300005f24 disk      connected  unconfigured unknown

```

```

c0::50020f2300006107 disk      connected  unconfigured unknown
c1                      fc-private connected  configured  unknown
c1::220203708b69c32b disk      connected  configured  unknown
c1::220203708ba7d832 disk      connected  configured  unknown
c1::220203708b8d45f2 disk      connected  configured  unknown
c1::220203708b9b20b2 disk      connected  configured  unknown

```

3 Configure the fabric device.

```
# cfgadm -c configure c0::50020f2300006077
```

4 Verify that the selected fabric device is configured.

```
# cfgadm -al
Ap_Id                Type      Receptacle Occupant  Condition
c0                   fc-fabric connected configured unknown
c0::50020f2300006077 disk      connected configured unknown
c0::50020f23000063a9 disk      connected unconfigured unknown
c0::50020f2300005f24 disk      connected unconfigured unknown
c0::50020f2300006107 disk      connected unconfigured unknown
c1                   fc-private connected configured unknown
c1::220203708b69c32b disk      connected configured unknown
c1::220203708ba7d832 disk      connected configured unknown
c1::220203708b8d45f2 disk      connected configured unknown
c1::220203708b9b20b2 disk      connected configured unknown

```

Notice that the Occupant column for both `c0` and `c0::50020f2300006077` displays as configured, indicating that the `c0` port has a configured occupant and that the `c0::50020f2300006077` device is configured.

5 Use the `show_SCSI_LUN` option to display FCP SCSI LUN information for multi-LUN SCSI devices.

The following code example shows that the physical devices connected through `Ap_Id` `c0::50020f2300006077` have four LUNs configured.

```
# cfgadm -al -o show_SCSI_LUN c0
Ap_Id                Type      Receptacle Occupant  Condition
c0                   fc-fabric connected configured unknown
c0::50020f2300006077,0 disk      connected configured unknown
c0::50020f2300006077,1 disk      connected configured unknown
c0::50020f2300006077,2 disk      connected configured unknown
c0::50020f2300006077,3 disk      connected configured unknown

```

The device is now available on the host using the Solaris Express Developer's Edition OS. The paths represent each SCSI LUN in the physical device represented by `c0::50020f2300006077`.

▼ To Configure Multiple Devices Without Multipathing

Make sure you first identify the devices visible to the host with the procedure “[Ensuring That LUN Level Information Is Visible](#)” on page 85. This procedure describes how to configure all unconfigured fabric devices that are attached to a fabric-connected host port. The port used as an example is `c0`.

- 1 Log in as root (`su - root`).
- 2 Identify the devices to be configured.

```
# cfmadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
<code>c0</code>	<code>fc-fabric</code>	<code>connected</code>	<code>unconfigured</code>	<code>unknown</code>
<code>c0::50020f2300006077</code>	<code>disk</code>	<code>connected</code>	<code>unconfigured</code>	<code>unknown</code>
<code>c0::50020f23000063a9</code>	<code>disk</code>	<code>connected</code>	<code>unconfigured</code>	<code>unknown</code>
<code>c0::50020f2300005f24</code>	<code>disk</code>	<code>connected</code>	<code>unconfigured</code>	<code>unknown</code>
<code>c0::50020f2300006107</code>	<code>disk</code>	<code>connected</code>	<code>unconfigured</code>	<code>unknown</code>
<code>c1</code>	<code>fc-private</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b69c32b</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708ba7d832</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b8d45f2</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b9b20b2</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>

- 3 Configure all of the unconfigured devices on the selected port.

```
# cfmadm -c configure c0
```

Note – This operation repeats the configure operation of an individual device for all the devices on `c0`. This can be time consuming if the number of devices on `c0` is large.

- 4 Verify that all devices on `c0` are configured.

```
# cfmadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
<code>c0</code>	<code>fc-fabric</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c0::50020f2300006077</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c0::50020f23000063a9</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c0::50020f2300005f24</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c0::50020f2300006107</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1</code>	<code>fc-private</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b69c32b</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708ba7d832</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b8d45f2</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>
<code>c1::220203708b9b20b2</code>	<code>disk</code>	<code>connected</code>	<code>configured</code>	<code>unknown</code>

5 Use the `show_SCSI_LUN` command to display FCP SCSI LUN information for multi-LUN SCSI devices.

The following code example shows that the physical devices represented by `c0::50020f2300006077` and `c0::50020f2300006107` each have four LUNs configured. The physical devices represented by `c0::50020f23000063a9` and `c0::50020f2300005f24` each have two LUNs configured.

```
# cfgadm -al -o show_SCSI_LUN c0
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077,0	disk	connected	configured	unknown
c0::50020f2300006077,1	disk	connected	configured	unknown
c0::50020f2300006077,2	disk	connected	configured	unknown
c0::50020f2300006077,3	disk	connected	configured	unknown
c0::50020f23000063a9,0	disk	connected	configured	unknown
c0::50020f23000063a9,1	disk	connected	configured	unknown
c0::50020f2300005f24,0	disk	connected	configured	unknown
c0::50020f2300005f24,1	disk	connected	configured	unknown
c0::50020f2300006107,0	disk	connected	configured	unknown
c0::50020f2300006107,1	disk	connected	configured	unknown
c0::50020f2300006107,2	disk	connected	configured	unknown
c0::50020f2300006107,3	disk	connected	configured	unknown

Configuring Device Nodes With Solaris Multipathing Software Enabled

This section describes how to perform fabric device configuration steps on a host that has the multipathing software enabled.

The devices attached to the fabric-connected host port are not configured by default and so are not available to the host using the Solaris Express Developer's Edition OS. Use the `cfgadm` `configure` and `cfgadm` `unconfigure` commands to manage device node creation for fabric devices. See the `cfgadm_fp(1M)` man page for additional information. The procedures in this section illustrate steps to detect fabric devices that are visible on a host and to configure them as multipathing devices to make them available to the host using the Solaris Express Developer's Edition OS.

The device information that you supply, and that is displayed by the `cfgadm(1M)` command, depends on your system configuration. For more information on the `cfgadm` command, see the `cfgadm_fp(1M)` and `cfgadm(1M)` man pages.

▼ To Configure Individual Devices With Sun StorageTek Traffic Manager Software

This sample procedure uses fabric-connected host ports `c0` and `c2` to configure fabric devices as multipath devices on a host that has the Sun StorageTek Traffic Manager software enabled.

The `cfgadm -c configure` command for fabric devices is the same regardless of whether multipathing is enabled.

- 1 Log in as root (`su - root`).
- 2 Identify the port WWN of the device to be configured as the Sun StorageTek Traffic Manager software device.

Look for devices on a fabric-connected host port marked as `fc-fabric`. These devices are the devices you can configure with the `cfgadm -c configure` command.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   unconfigured unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c1             fc-private   connected   configured   unknown
c1::220203708b69c32b disk         connected   configured   unknown
c1::220203708ba7d832 disk         connected   configured   unknown
c1::220203708b8d45f2 disk         connected   configured   unknown
c1::220203708b9b20b2 disk         connected   configured   unknown
c2             fc-fabric    connected   unconfigured unknown
c2::50020f2300005f24 disk         connected   unconfigured unknown
c2::50020f2300006107 disk         connected   unconfigured unknown
```

In the above example, the `c0::50020f2300006077` and `c2::50020f2300006107` `Ap_Ids` represent the same storage device with different port WWNs for the storage device controllers. The `c0` and `c2` host ports are enabled for use by the Sun StorageTek Traffic Manager software.

- 3 Configure the fabric device and make the devices available to the host.

```
# cfgadm -c configure c0::50020f2300006077 c2::50020f2300006107
```

- 4 Verify that the selected devices are configured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured   unknown
c0::50020f2300006077 disk         connected   configured   unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c1             fc-private   connected   configured   unknown
c1::220203708b69c32b disk         connected   configured   unknown
c1::220203708ba7d832 disk         connected   configured   unknown
```

```

c1::220203708b8d45f2 disk      connected  configured  unknown
c1::220203708b9b20b2 disk      connected  configured  unknown
c2                               fc-fabric  connected  configured  unknown
c2::50020f2300005f24 disk      connected  unconfigured unknown
c2::50020f2300006107 disk      connected  configured  unknown

```

Notice that the Occupant column of `c0` and `c0::50020f2300006077` specifies `configured`, which indicates that the `c0` port has at least one configured occupant and that the `c0::50020f2300006077` device is configured. The same change has been made in `c2` and `c2::50020f2300006107`.

When the configure operation has been completed without an error, multipathing-enabled devices are created on the host. If the physical device represented by `c0::50020f2300006077` and `c2::50020f2300006107` has multiple SCSI LUNs configured, each LUN is configured as a Sun StorageTek Traffic Manager software device. The example below shows that two LUNs are configured through `c0::50020f2300006077` and `c2::50020f2300006107`. Each `Ap_Id` is associated with a path to those Sun StorageTek Traffic Manager software devices.

```

# cfmadm -al -o show_SCSI_LUN c0::50020f2300006077\ c2::50020f2300006107
Ap_Id      Type      Receptacle  Occupant  Condition
c0::50020f2300006077,0 disk      connected   configured unknown
c0::50020f2300006077,1 disk      connected   configured unknown
c2::50020f2300006107,0 disk      connected   configured unknown
c2::50020f2300006107,1 disk      connected   configured unknown

```

The example above shows that the following two multipathing-enabled devices are created for the device represented by `c0::50020f2300006077` and `c2::50020f2300006107`:

▼ To Configure Multiple Devices With Multipathing

Before you configure or remove device nodes, be sure to first identify the fabric devices by using the procedure “[Ensuring That LUN Level Information Is Visible](#)” on page 85.

In this example, an `Ap_Id` on a fabric-connected host port is a path to the Sun StorageTek Traffic Manager software device. For example, all devices with a path through `c2` are to be configured, but none through `c0` are to be configured. `c2` is an attachment point from the host to the fabric, whereas `c2::50020f2300006107` is an attachment point from the storage to the fabric. A host detects all the storage devices in a fabric for which it is configured.

Configuring an `Ap_Id` on the device that has already been configured through another `Ap_Id` results in an additional path to the previously configured device. A new device node is not created in this case. The device node is created only the first time an `Ap_Id` to the corresponding device is configured.

1 Log in as root (su - root)..

2 Identify the fabric-connected host port to be configured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured  unknown
c0::50020f2300006077 disk         connected   configured  unknown
c0::50020f23000063a9 disk         connected   configured  unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
c2             fc-fabric    connected   unconfigured unknown
c2::50020f2300005f24 disk         connected   unconfigured unknown
c2::50020f2300006107 disk         connected   unconfigured unknown
```

Devices represented by Ap_Ids `c0::50020f2300006077` and `c2::50020f2300006107` are two paths to the same physical device, with `c0::50020f2300006077` already configured. Configure the unconfigured devices on the selected port. This operation repeats the `configure` command of an individual device for all the devices on `c2`. This can be time-consuming if the number of devices on `c2` is large.

```
# cfgadm -c configure c2
```

3 Verify that all devices on c2 are configured.

```
# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c0             fc-fabric    connected   configured  unknown
c0::50020f2300006077 disk         connected   configured  unknown
c0::50020f23000063a9 disk         connected   configured  unknown
c1             fc-private   connected   configured  unknown
c1::220203708b69c32b disk         connected   configured  unknown
c1::220203708ba7d832 disk         connected   configured  unknown
c1::220203708b8d45f2 disk         connected   configured  unknown
c1::220203708b9b20b2 disk         connected   configured  unknown
c2             fc-fabric    connected   configured  unknown
c2::50020f2300005f24 disk         connected   configured  unknown
c2::50020f2300006107 disk         connected   configured  unknown
```

Notice that the `Occupant` column of `c2` and all of the devices under `c2` is marked as `configured`.

The `show_SCSI_LUN` command displays FCP SCSI LUN information for multiple LUN SCSI devices. The following code example shows that the physical devices connected through by `c2::50020f2300006107` and `c2::50020f2300005f24` each have two LUNs configured.

```
# cfgadm -al -o show_SCSI_LUN c2
Ap_Id          Type          Receptacle  Occupant    Condition
c2             fc-fabric    connected   configured  unknown
c2::50020f2300005f24,0 disk         connected   configured  unknown
```

```

c2::50020f2300005f24,1 disk    connected    configured    unknown
c2::50020f2300006107,0 disk    connected    configured    unknown
c2::50020f2300006107,1 disk    connected    configured    unknown

```

Unconfiguring Fabric Devices

This section provides information about unconfiguring the Sun StorageTek Traffic Manager software.

Unconfiguring a Fabric Device

Before you unconfigure a fabric device, stop all activity to the device and unmount any file systems on the fabric device. See the administration documentation for the Solaris Express Developer's Edition OS for unmounting instructions. If the device is under any volume manager's control, see the documentation for your volume manager before unconfiguring the device.

▼ To Unconfigure a Fabric Device

This procedure describes how to unconfigure a fabric device that is attached to the fabric-connected host port c0.

1 Log in as root (su - root)..

2 Identify the device to be unconfigured.

Only devices on a fabric-connected host port can be unconfigured.

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077	disk	connected	configured	unknown
c0::50020f23000063a9	disk	connected	configured	unknown
c1	fc-private	connected	configured	unknown
c1::220203708b69c32b	disk	connected	configured	unknown
c1::220203708ba7d832	disk	connected	configured	unknown

3 Unconfigure the fabric device.

```
# cfgadm -c unconfigure c0::50020f2300006077
```

4 Verify that the selected fabric device is unconfigured.

```
# cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown

```

c0::50020f2300006077 disk      connected  unconfigured unknown
c0::50020f23000063a9 disk      connected  configured   unknown
c1                fc-private connected  configured   unknown
c1::220203708b69c32b disk      connected  configured   unknown
c1::220203708ba7d832 disk      connected  configured   unknown

```

▼ To Unconfigure All Fabric Devices on a Fabric-Connected Host Port

This procedure describes how to unconfigure all configured fabric devices that are attached to a fabric-connected host port.

1 Log in as root (su - root)..

2 Identify the fabric devices to be unconfigured.

Only devices on a fabric-connected host port can be unconfigured.

```

# cfgadm -aL
Ap_Id          Type          Receptacle  Occupant  Condition
c0              fc-fabric     connected   configured unknown
c0::50020f2300006077 disk         connected   configured unknown
c0::50020f23000063a9 disk         connected   configured unknown
c1              fc-private    connected   configured unknown
c1::220203708b69c32b disk         connected   configured unknown
c1::220203708ba7d832 disk         connected   configured unknown

```

3 Stop all activity to each fabric device on the selected port and unmount any file systems on each fabric device.

If the device is under any volume manager's control, see the documentation for your volume manager before unconfiguring the device.

```
# cfgadm -c unconfigure c0
```

4 Unconfigure all of the configured fabric devices on a selected port.

Note – This operation repeats the unconfigure operation of an individual device for all the devices on c0. This process can be time-consuming if the number of devices on c0 is large.

5 Verify that all the devices on c0 are unconfigured.

```

# cfgadm -aL
Ap_Id          Type          Receptacle  Occupant  Condition
c0              fc-fabric     connected   unconfigured unknown
c0::50020f2300006077 disk         connected   unconfigured unknown
c0::50020f23000063a9 disk         connected   unconfigured unknown
c1              fc-private    connected   configured   unknown
c1::220203708b69c32b disk         connected   configured   unknown
c1::220203708ba7d832 disk         connected   configured   unknown

```

Notice that the Occupant column of `c0` and all the fabric devices attached to it are displayed as unconfigured.

▼ To Unconfigure a Fabric Device Associated With Sun StorageTek Traffic Manager Multipathing-Enabled Devices

This procedure shows fabric-connected host ports `c0` and `c2` to illustrate how to unconfigure fabric devices associated with Sun StorageTek Traffic Manager software devices.

- 1 Log in as root (`su - root`).
- 2 Identify the port WWN of the fabric device to be unconfigured.

```
# cfgadm -aL
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected   configured unknown
c0::50020f2300006077 disk      connected   configured unknown
c0::50020f23000063a9 disk      connected   configured unknown
c1                   fc-private connected   configured unknown
c1::220203708b69c32b disk      connected   configured unknown
c1::220203708ba7d832 disk      connected   configured unknown
c2                   fc-fabric connected   configured unknown
c2::50020f2300005f24 disk      connected   configured unknown
c2::50020f2300006107 disk      connected   configured unknown
```

In this example, the `c0::50020f2300006077` and `c2::50020f2300006107` Ap_Ids represent different port WWNs for the same device associated with a Sun StorageTek Traffic Manager software device. The `c0` and `c2` host ports are enabled for use by the Sun StorageTek Traffic Manager software.

- 3 Stop all device activity to each fabric device on the selected port and unmount any file systems on each fabric device.

If the device is under any volume manager's control, see the documentation for your volume manager for maintaining the fabric device.

- 4 Unconfigure fabric devices associated with the device.

Only devices on a fabric-connected host port can be unconfigured through the `cfgadm -c unconfigure` command.

```
# cfgadm -c unconfigure c0::50020f2300006077 c2::50020f2300006107
```

Note – You can remove a device from up to eight paths individually, as in the example command `cfgadm -c unconfigure c0::1111, c1::2222, c3::3333`, and so on. As an alternative, you can remove an entire set of paths from the host, as in the example `cfgadm -c unconfigure c0`.

5 Verify that the selected devices are unconfigured.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected   configured unknown
c0::50020f2300006077 disk      connected   unconfigured unknown
c0::50020f23000063a9 disk      connected   configured  unknown
c1                   fc-private connected   configured  unknown
c1::220203708b69c32b disk      connected   configured  unknown
c1::220203708ba7d832 disk      connected   configured  unknown
c2                   fc-fabric connected   configured  unknown
c2::50020f2300005f24 disk      connected   configured  unknown
c2::50020f2300006107 disk      connected   unconfigured unknown
```

Notice that the Ap_Ids `c0::50020f2300006077` and `c2::50020f2300006107` are unconfigured. The Occupant column of `c0` and `c2` still displays those ports as configured because they have other configured occupants.

The Sun StorageTek Traffic Manager software devices associated with the Ap_Ids `c0::50020f2300006077` and `c2::50020f2300006107` are no longer available to the host using the Solaris Express Developer's Edition OS. The following two devices are removed from the host:

```
/dev/rdisk/c6t60020F20000061073AC8B52D000B74A3d0s2
```

```
/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
```

▼ To Unconfigure One Path to a Multipathing Device

In contrast to the procedure in the preceding section, this procedure shows how to unconfigure one device associated with `c2::50020f2300006107` and leave the other device, `50020f2300006077`, configured. Only devices on a fabric-connected host port can be unconfigured through the `cfgadm unconfigure` command.

1 Log in as root (su - root)..

2 Identify the Ap_Id of the Sun StorageTek Traffic Manager software device to be unconfigured.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected   configured  unknown
c0::50020f2300006077 disk      connected   configured  unknown
c0::50020f23000063a9 disk      connected   configured  unknown
c1                   fc-private connected   configured  unknown
c1::220203708b69c32b disk      connected   configured  unknown
c1::220203708ba7d832 disk      connected   configured  unknown
c2                   fc-fabric connected   configured  unknown
c2::50020f2300005f24 disk      connected   configured  unknown
c2::50020f2300006107 disk      connected   configured  unknown
```


In this example, `c0::50020f2300006077` and `c2::50020f2300006107` Ap_Ids represent different port WWNs for the same device.

3 Unconfigure the Ap_Id associated with the device.

Note – If the Ap_Id represents the last configured path to the device, stop all activity to the path and unmount any file systems on it. If the multipathing device is under any volume manager's control, see the documentation for your volume manager for maintaining the fabric device.

In the example that follows, the path represented as `c2::50020f2300006107` is unconfigured, and `c0::50020f2300006077` remains configured to show how you can unconfigure just one of multiple paths for a multipathing device.

```
# cfgadm -c unconfigure c2::50020f2300006107
```

4 Verify that the selected path `c2::50020f2300006107` is unconfigured.

```
# cfgadm-al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077	disk	connected	configured	unknown
c0::50020f23000063a9	disk	connected	configured	unknown
c1	fc-private	connected	configured	unknown
c1::220203708b69c32b	disk	connected	configured	unknown
c1::220203708ba7d832	disk	connected	configured	unknown
c2	fc-fabric	connected	configured	unknown
c2::50020f2300005f24	disk	connected	configured	unknown
c2::50020f2300006107	disk	connected	unconfigured	unknown

The devices associated with that Ap_Id are still available to a host using the Solaris Express Developer's Edition OS through the other path, represented by `c0::50020f2300006077`. A device can be connected to multiple Ap_Ids and an Ap_Id can be connected to multiple devices.

```
/dev/rdisk/c6t60020F20000061073AC8B52D000B74A3d0s2
and
/dev/rdisk/c6t60020F20000061073AC8B4C50004ED3Ad0s2
```

▼ To Unconfigure All Fabric-Connected Devices With Multipathing Enabled

An Ap_Id on a fabric-connected host port is a path to the Sun StorageTek Traffic Manager software device.

When a Sun StorageTek Traffic Manager software device has multiple Ap_Ids connected to it, the device is still available to the host using the Solaris Express Developer's Edition OS after you unconfigure an Ap_Id. After you unconfigure the last Ap_Id, no additional paths remain and

the device is unavailable to the host using the Solaris Express Developer's Edition OS. Only devices on a fabric-connected host port can be unconfigured.

1 Log in as root (su - root)..

2 Identify the devices to be unconfigured.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected   configured unknown
c0::50020f2300006077 disk      connected   configured unknown
c0::50020f23000063a9 disk      connected   configured unknown
c1                   fc-private connected   configured unknown
c1::220203708b69c32b disk      connected   configured unknown
c1::220203708b9b20b2 disk      connected   configured unknown
c2                   fc-fabric connected   configured unknown
c2::50020f2300005f24 disk      connected   configured unknown
```

3 Unconfigure all of the configured devices on the selected port.

```
# cfgadm -c unconfigure c2
```

Note – This operation repeats the unconfigure command of an individual device for all devices on c2. This process can be time-consuming if the number of devices on c2 is large.

4 Verify that all devices on c2 are unconfigured.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c0                   fc-fabric connected   configured unknown
c0::50020f2300006077 disk      connected   configured unknown
c1                   fc-private connected   configured unknown
c1::220203708b69c32b disk      connected   configured unknown
c1::220203708ba7d832 disk      connected   configured unknown
c2                   fc-fabric connected   unconfigured unknown
c2::50020f2300005f24 disk      connected   unconfigured unknown
c2::50020f2300006107 disk      connected   unconfigured unknown
```

Notice that the Occupant column lists c2 and all the devices attached to c2 as unconfigured.

Supported FC-HBA API

This appendix contains the list of supported and unsupported FC-HBA Interfaces. For further information regarding the API, refer to the Sun StorageTek Traffic Manager multipathing [“Software Features” on page 12](#).

Supported Fibre Channel HBA API

TABLE B-1 Supported and Unsupported FC-HBA Interfaces

FC - HBA Interfaces	Supported or Not Supported by Sun StorageTek Traffic Manager Software
HBA_GetVersion	Supported
HBA_LoadLibrary	Supported
HBA_FreeLibrary	Supported
HBA_GetNumberOfAdapters	Supported
HBA_GetAdapterName	Supported
HBA_OpenAdapter	Supported
HBA_CloseAdapter	Supported
HBA_GetAdapterAttributes	Supported
HBA_GetAdapterPortAttributes	Supported
HBA_GetDiscoveredPortAttributes	Supported
HBA_GetPortAttributesbyWWN	Supported
HBA_SendCTPassThru	Supported

TABLE B-1 Supported and Unsupported FC-HBA Interfaces (Continued)

FC - HBA Interfaces	Supported or Not Supported by Sun StorageTek Traffic Manager Software
HBA_SendCTPassThruV2	Supported
HBA_RefreshInformation	Supported
HBA_GetFcpTargetMapping	Supported
HBA_SendScsiInquiry	Supported
HBA_SendReportLuns	Supported
HBA_SendReadCapacity	Supported
HBA_GetPortStatistics	Supported
HBA_ResetStatistics	Not supported
HBA_GetFcpPersistentBinding	Not supported
HBA_GetEventBuffer	Not supported
HBA_SetRNIDMgmtInfo	Supported
HBA_GetRNIDMgmtInfo	Supported
HBA_SendRNID	Supported
HBA_SendRNIDV2	Supported
HBA_ScsiInquiryV2	Supported
HBA_ScsiReportLUNsv2	Supported
HBA_ScsiReadCapacityV2	Supported
HBA_OpenAdapterByWWN	Supported
HBA_RefreshAdapterConfiguration	Supported
HBA_GetVendorLibraryAttributes	Supported
HBA_GetFcpTargetMappingV2	Supported
HBA_SendRPL	Not supported
HBA_SendRPS	Not supported
HBA_SendSRL	Not supported
HBA_SendLIRR	Not supported
HBA_SendRLS	Supported
HBA_RemoveCallback	Supported

TABLE B-1 Supported and Unsupported FC-HBA Interfaces *(Continued)*

FC - HBA Interfaces	Supported or Not Supported by Sun StorageTek Traffic Manager Software
HBA_RegisterForAdapterEvents	Supported
HBA_RegisterForAdapterAddEvents	Supported
HBA_RegisterForAdapterPortEvents	Supported
HBA_RegisterForAdapterPortStatEvents	Not supported
HBA_RegisterForTargetEvents	Supported
HBA_RegisterForAdapterLinkEvents	Not supported
HBA_RegisterForAdapterTargetEvents	Supported
HBA_GetFC4Statistics	Not supported
HBA_GetFCPStatistics	Not supported
HBA_GetBindingCapability	Not supported
HBA_GetBindingSupport	Not supported
HBA_SetBindingSupport	Not supported
HBA_SetPersistentBindingV2	Not supported
HBA_GetPersistentBindingV2	Not supported
HBA_RemovePersistentBinding	Not supported
HBA_RemoveAllPersistentBindings	Not supported

Multipathing Troubleshooting

This appendix provides solutions to potential problems that might occur while running the Sun StorageTek Traffic Manager software.

This appendix contains the following sections:

- “System Fails to Boot During `stmsboot`” on page 103
- “System Crashes During `stmsboot`” on page 104

System Fails to Boot During `stmsboot`

During the reboot following `stmsboot enable (-e)`, `stmsboot disable (-d)`, or `stmsboot update (-u)` operations if the system fails to boot but lets you enter single user mode, perform the following to recover.

▼ To Recover From Single User Mode

- 1 Type the root password and enter single user mode.
- 2 Restart the `mpxio-upgrade` service.

```
# svcadm restart svc:/system/device/mpxio-upgrade:default
```

If this command is not successful, follow the instructions in the next section to recover your original configuration.

Note – In prior versions of the Solaris Express Developer's Edition OS, this service was known as `svc:/platform/sun4u/mpxio-upgrade:default`.

System Crashes During `stmsboot`

During the reboot following `stmsboot enable (-e)`, `stmsboot disable (-d)`, or `stmsboot update (-u)` operations, if your system hangs, panics, or fails to boot, follow the instructions below to recover your original configuration.

▼ To Recover from a System Crash

1 Boot the system from another disk, CD, DVD, or over the network.

2 Run the `fsck your-root-device` command where *your-root-device* is your original root device.

```
# fsck your-root-device
```

3 Mount your original root device.

```
# mount your-root-device /mnt
```

4 Restore your original `fp.conf` file (for FC multipathing) or `mpt.conf` (for SAS multipathing) as follows.

▪ **If you ran the `stmsboot -e` command or `stmsboot -d` command:**

- For SAS multipathing, type the command `cp /mnt/etc/mpxio/mpt.conf /mnt/kernel/drv/mpt.conf`
- For FC multipathing, type the command `cp /mnt/etc/mpxio/fp.conf /mnt/kernel/drv/fp.conf`

▪ **If you ran the `stmsboot -u` command, and you modified either the `fp.conf` file or the `mpt.conf` file, undo the changes you made to this file by editing either the `/mnt/kernel/drv/fp.conf` or the `/mnt/kernel/drv/mpt.conf` files.**

5 Undo any other changes you made to the multipathing configuration prior to running the `stmsboot` command.

For example, if you modified the `scsi_vhci.conf` file, undo the changes you made to this file by editing the `/mnt/kernel/drv/scsi_vhci.conf` file.

If you modified the device settings of your storage arrays, restore their original settings.

6 Restore your original `/etc/vfstab` file:

```
# cp /mnt/etc/mpxio/vfstab /mnt/etc/vfstab
```

The `/mnt/etc/mpxio/vfstab` file is a copy your original `/etc/vfstab` file that the `stmsboot` command saved prior to updating your `vfstab` file. A `/mnt/etc/mpxio/vfstab` file will not exist if the `stmsboot` command has not modified your `vfstab` file.

- 7 If the system is running on a Solaris Express Developer's Edition OS x86 based system, perform the following steps:
 - a. Restore your original `/boot/solaris/bootenv.rc` file.

```
# cp /mnt/etc/mpxio/bootenv.rc /mnt/boot/solaris/bootenv.rc
```

The `/mnt/etc/mpxio/bootenv.rc` file is a copy your original `/boot/solaris/bootenv.rc` file that the `stmsboot` command saved prior to updating your `bootenv.rc` file. A `/mnt/etc/mpxio/bootenv.rc` file will not exist if the `stmsboot` command has not modified your `bootenv.rc` file.
 - b. Update the boot archive.

```
# bootadm update-archive -R /mnt
```
- 8 Disable the `mpxio-upgrade` service:

```
# /usr/sbin/svccfg -f /mnt/etc/mpxio/svccfg_recover
```
- 9 Type `sync` to synchronize the file systems.
- 10 Unmount the device.

```
# umount /mnt
```
- 11 Reboot the system.

Index

Numbers and Symbols

24-bit FC addressing devices, 84

A

ANSI standard Multipath Management API, 21-35

Ap_Id, 85

automatic failback, configuration, 47

B

boot disk, 51

Booting, 13

broadcasting, 63

C

cfgadm, 50

-c, 87

-l, 85, 86

cfgadm -al -o show_SCSI_LUN, 85

cfgadm -c unconfigure, 95

cfgadm(1M), 85

cfgadm configure, 86

cfgadm_fp(1M), 89

configuration

device, 85

dump, 38

fabric devices, 86

FCoE, 58-60

configuration (*Continued*)

Fibre Channel overview, 17-19

manual, 83-84

multipathing, 38

SAS devices, 62

tasks, 17

configure

automatic failback, 47

fabric-connected devices, 49-51

multipathing feature, 37, 49-51

NPIV ports, 53-56

SAN devices, 49-51

SAS devices, 62

third-party devices, 43-44, 44

virtual Fibre Channel ports, 53-56

considerations

device name changes, 38

device specific, 38

fabric boot devices, 51

IPFC, 63-66

multipathing, 38

multipathing, enabling and disabling, 38-41

per-port configuration, 41-43

SAN devices, 49-50

SAS devices, 61

stmsboot command, 18

StorageTek Traffic Manager, 18-19

third-party device configuration, 43-44

D

- device
 - asymmetrical, 14
 - configuration, 86
 - storage, 14
- device management, 15
- devices, tape, 12
- disabling, multipathing, 38-41
- discovery
 - dynamic, 61
 - SAS devices, 61

E

- enabling, multipathing, 38-41

F

- fabric-connected host ports, 85, 90
- fabric device configuration, 86
- fc-fabric, 90
- FCAL, 12
- FCoE
 - configuring ports, 58-60
 - creating ports, 58-59
 - deleting ports, 59
 - limitations, 57
 - reporting status on ports, 59-60
- Fibre Channel
 - arbitrated loop, 12
 - booting Solaris 10 x86/x64, 69-78
 - configuration overview, 17-19
 - configuring IPFC SAN devices, 63-68
 - determining adapter port instances, 64-66
 - enabling and disabling multipathing, 38-41
 - HBA, 13
 - IPFC considerations, 63-66
 - listing LUN level information, 85-86
 - multipathing configuration, 37
 - supported HBA interfaces, 99-101
- Fibre Channel ports, configuring, 53-56

I

- initiator ports, displaying properties, 23
- Internet Protocol over Fibre Channel, *See* IPFC
- IPFC, considerations, 63-66

L

- limitations, FCoE, 57
- logical units, listing, 23-24
- LUN
 - masking, 49
 - recognition, 51
 - SCSI devices, 92
- luxadm(1M), 47

M

- modinfo, 85
- mpathadm commands, 21-35
 - automatic failback, configuring, 28
 - configure automatic failback, 28
 - disable a path, 34-35
 - enable a path, 32-33
 - initiator ports
 - discovered, 23
 - properties, 23
 - logical unit with a particular name, 27-28
 - logical units
 - associated with target port, 25
 - discovered, 23
 - fail over, 30-32
 - properties and associated path, 24
- mpt driver, 61
- Multipath Management API plug-in
 - library, 21
 - properties, 21
- multipathing
 - boot considerations, 51
 - configuring automatic failback, 47
 - configuring device nodes, 89-93
 - configuring third-party storage devices, 43-46
 - disabled, device node configuration, 86
 - disabled, multiple device configuration, 88-89

multipathing (*Continued*)

- enabling and disabling, 38-41
- enabling and disabling on a per-port basis, 41-43
- fabric device node configuration, 84
- features, 14-15
- manual configuration, 83-84
- mpathadm command, 21-35
- SAS boot device configuration, 62
- SAS considerations, 61
- support, 21-35
- troubleshooting, 103-105

N

- NPIV ports, configuring, 53-56

P

- per-port configuration
 - considerations, 41-43
 - enabling and disabling multipathing, 41-43
- persistent binding of tape devices, 79-81
- physical device, 87
- Promiscuous mode, 63

R

- REPORT LUNS command, 43

S

- SAS devices
 - configuration, 62
 - dynamic discovery, 61
 - multipathing considerations, 61
- sd driver, 61
- Serial Attached SCSI, *See* SAS
- show_SCSI_LUN, 87, 92
- snoop(1M), 63
- Solaris 10 x86/x64 FC OS booting procedure, 69-78
- ssd driver, 85

- st driver, 85
- stmsboot command, 38
 - boot considerations, 51
 - considerations, 18
 - determining device paths, 38
 - enabling and disabling multipathing, 15
 - SAN device considerations, 50
 - system crashes, 104-105
- storage device, 47

T

- T10 standard, 21-35
- T11 standard, 13
- tape device persistent binding, 79-81
- To Detect Fabric Devices Visible on a Host, 91
- troubleshooting, multipathing, 103-105

