

FC Solaris® Host Utilities 5.0 for Native OS with MPxIO and Veritas™ Dynamic Multipathing **Installation and Setup Guide**

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 215-04282_A0
December 2008

Copyright and trademark information

Copyright information

Copyright © 1994-2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

The third-party copyright statements and licenses that apply to open source software components distributed with the product are contained in the section “Third-Party Copyrights and Licenses.”

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, VFM (Virtual File Manager), and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by

Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Host Utilities Installation and Setup Overview	1
	Host Utilities Overview.	2
	Supported Solaris, Data ONTAP features, configurations	5
	Information on the stacks supported by the Host Utilities	8
	FC Solaris Host Utilities installation overview.	11
	Where to go for more information	15
Chapter 2	Veritas DMP: Managing the Driver Software.	19
	Before you begin working with the drivers.	20
	Working with the Emulex LPFC software	24
	Working with the Sun Drivers for Emulex HBAs (emlxs).	34
	Working with the Sun Native Drivers for QLogic HBAs (qlc)	37
Chapter 3	Installing the Host Utilities	41
	Installing the Host Utilities software package	42
	Upgrading the FC Solaris Host Utilities	48
	Removing an FC Solaris Host Utilities or Attach Kit installation	49
Chapter 4	Veritas DMP: Configuring the Solaris Host.	53
	Completing the installation and configuring system parameters.	54
	Overview of using the basic_config command.	56
	Configuring parameters on systems using LPFC drivers	57
	Configuring parameters on systems using Sun native drivers	62
Chapter 5	MPxIO: Configuring the Solaris Host	65
	Completing the installation and configuring system parameters.	66

Chapter 6	Veritas DMP: Installing the ASL and APM.	71
	Overview of the Array Support Library and Array Policy Module	72
	Installing or Upgrading the ASL and the APM.	74
	Example of uninstalling and installing the ASL and APM.	79
	Obtaining multipathing information	83
	Optimizing the Veritas DMP Restore Daemon.	88
	Understanding ASL error message types.	90
Chapter 7	Basic Tasks for Working with LUNs.	91
	Overview of LUN configuration and management.	92
	Creating and mapping igroups and LUNs	96
Chapter 8	Veritas DMP: Configuring and Managing LUNs.	99
	(LPFC drivers) Creating persistent bindings	100
	(LPFC drivers) Configuring the host to discover LUNs	102
	(Sun native drivers) Configuring the host to discover LUNs.	105
	Labeling the LUN and configuring Volume Management software	106
Chapter 9	MPxIO: Configuring and Managing LUNs	109
	Configuring the host to discover LUNs	110
	Completing the configuration.	111
Chapter 10	Using sanlun	115
	Displaying host LUN information	116
	Displaying information about the storage system paths	119
	Displaying host HBA information	124
Chapter 11	Tips and Troubleshooting.	127
	Confirming that your system is set up correctly	128
	Using Diagnostic Utilities	140

	Collecting storage system information	141
	Collecting Solaris information	143
	Collecting switch information	144
	Known issues with the Host Utilities.	146
Appendix A	Veritas DMP: Using the create_binding.pl Script.	149
	The create_binding.pl script	150
	Before you begin	152
	Using the create_binding.pl script	154
Appendix B	Veritas DMP: Configuring SAN Boot	159
	Configuration overview.	160
	(Sun native drivers) Setting up the HBA	162
	(LPFC drivers) Setting up the HBA with Emulex drivers	168
	Creating the bootable LUN.	172
	Performing a Direct Installation to a SAN Boot LUN	178
	(LPFC, Sun Native Drivers) Copy bootblks, data to SAN boot LUN	179
	Modifying OpenBoot.	183
Appendix C	MPxIO: Configuring SAN Boot	191
	Configuration overview.	192
	Performing a direct installation.	194
	File system dump, restore for a SPARC system without MPxIO	195
	File system dump, restore for an x86/x64 system with MPxIO	200
	File system dump, restore for an x86/x64 system without MPxIO.	204
Appendix D	Third-Party Copyrights and Licenses	207
	Index	233

Preface

About this guide

This guide describes how to configure a Solaris® host in a SAN environment. The host accesses data on a NetApp® storage system that runs Data ONTAP® software with the Fibre Channel Protocol (FCP).

See the NetApp Interoperability Matrix, available on the NetApp on the Web™ (NOW) site at <http://now.netapp.com/matrix/mtx/login.do/> for information on the specific Solaris versions supported.

Audience

This guide is for system installers and administrators who:

- ◆ Are familiar with the Sun® Solaris operating system
- ◆ Understand how to configure and administer a storage system
- ◆ Understand how to use FC to store and retrieve data on a storage system.

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (controllers, filers, FAS appliances, and NearStore systems) are all storage systems and are also sometimes called *filers* or *storage appliances*.
- ◆ *igroup* refers to a collection of unique identifiers, either FC WWPNs in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs.
- ◆ *LUN* (Logical Unit Number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Initiator* refers to the system component that originates an I/O command over an I/O bus or network.
- ◆ *Target* refers to the storage system that receives an FC I/O request and serves the data requested by the initiator.

Command conventions

You can enter storage system commands on the system console or from any client that can obtain access to the storage system using a Telnet session.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcstest adaptername</code> command, you enter the characters “fcstest” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter, enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Also used to mean pressing one or more keys and then pressing the Enter key.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.

Special messages

This guide contains special messages that are described as follows:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

How to send your comments

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this publication, you can send us comments electronically to doccomments@netapp.com. To help us direct the comments to the correct division, please put FC Solaris in the subject line.

About this chapter

This chapter provides an introduction to the FC Solaris Host Utilities for Native OS (MPxIO) and Veritas.

In addition, this chapter provides a high-level overview of the tasks you need to perform to install the Host Utilities and set up your system. It also includes pointers to other documentation that might be helpful in working with your Host Utilities setup.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Host Utilities Overview”](#) on page 2
- ◆ [“Supported Solaris, Data ONTAP features, configurations”](#) on page 5
- ◆ [“Information on the stacks supported by the Host Utilities”](#) on page 8
- ◆ [“FC Solaris Host Utilities installation overview”](#) on page 11
- ◆ [“Where to go for more information”](#) on page 15

Host Utilities Overview

Important information about Host Utilities

The Host Utilities provide components that enable users to connect Solaris hosts to NetApp storage systems. They support different Solaris environments:

- ◆ Veritas. This environment uses Veritas Storage Foundation to provide volume management and multipathing. For multipathing, it uses Veritas Dynamic Multipathing (DMP) with either Emulex® LightPulse® Fibre Channel (LPFC) driver or Sun native drivers (Leadville).
- ◆ Native OS. This environment uses features of the Solaris operating system. It uses Sun™ StorEdge™ SAN Foundation Software with Sun StorageTek™ Traffic Manager (MPxIO) and either the Sun Volume Manager, ZFS™ file system, or the Veritas Volume Manager.

Note

Going forward, this document uses the term MPxIO to refer to the Native OS environment.

To determine which environments and features the Host Utilities support, check the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Note

Previous versions of the Host Utilities were called FCP Solaris Attach Kits.

The tools contained in the Host Utilities apply to numerous components. Sometimes the same tools are used by Host Utilities for other operating systems. As a result, the tools do not always have the same version number as the Host Utilities version number. This is not a problem. The tools contained in the Host Utilities installation files are correct.

Overview of the FC Solaris Host Utilities

The Host Utilities consist of the following tools and information:

- ◆ The SAN Toolkit, which contains:
 - ❖ Configuration tools that help you configure host, HBA, system files, and create persistent bindings.You can configure the Solaris host using the tools provided in the FC Solaris Host Utilities, HBA-provided tools, or a combination FC Solaris Host Utilities and HBA-provided tools.

- ❖ The `san lun` utility that helps you manage LUNs and the host HBA.
- ❖ Diagnostic scripts that provide information about components in your configuration. Customer Support may ask for output from these scripts.
- ◆ Documentation that describes how to configure the Solaris host. It includes this setup guide, *Release Notes*, the *Host Settings Changed by the Host Utilities* document, and a *Quick Command Reference for FC Solaris Host Utilities for Native OS and Veritas*.

Note

The *Release Notes* are updated whenever new information relating to the Host Utilities is found. You should check the *Release Notes* before installing the Host Utilities to see if there are new details or corrections that apply to the installation.

NetApp qualifies new components on an ongoing basis. The NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do> contains the most current information on the supported configurations.

Software tools provided by the Host Utilities

The FC Solaris Host Utilities bundles the software tools into a SAN Toolkit.

Note

This toolkit is common across all the Solaris Host Utilities configurations: MPxIO and Veritas DMP as well as FC and iSCSI. As a result, some of its contents apply to one configuration, but not another. Having a program or file that does not apply to your configuration does not affect performance.

The toolkit contains the following components:

- ◆ `basic_config` command line tool. This tool modifies the Fibre Channel tunable variables to set them to values required by the FC driver. It provides the HBA time-out settings for both Sun and Veritas LPFC driver configurations.
- ◆ `create_binding.pl` command line tool. This tool enables you to bind the WWPNs of a single storage system's active ports to a target ID. It saves these bindings in the `/kernel/drv/lpfc.conf` file.
- ◆ `filer_info` script. This script collects information about the storage system.
- ◆ `mpxio_set` script. This script is only necessary if you had installed the Host Utilities 3.1 or earlier. This script removes the symmetric options from the `scsi_vhci.conf` file.
- ◆ `san_version` command. This command displays the version of the SAN Toolkit that you are running.

- ◆ `sanlun` utility. This utility displays information about LUNs on the storage system that are available to this host.
- ◆ `ssd_config.pl` script. The `basic_config` tool calls this script to make setting changes for the FC and SCSI drivers. You should never run this script.
- ◆ `solaris_info` script. This script collects configuration information about your host.
- ◆ Additional diagnostic scripts. The following scripts can also be used to collect troubleshooting information:
 - ❖ `brocade_info` collects information about Brocade switches installed in the network.
 - ❖ `cisco_info` collects information about Cisco switches installed in the network.
 - ❖ `mcdata_info` collects information about McData switches installed in the network.
 - ❖ `qlogic_info` collects information about QLogic switches installed in the network.
 - ❖ `SHsupport.pm` is the support library for the `*info` scripts. You should never run this script.
 - ❖ `Telnet.pm` support library for `*info` scripts. You should never run this script.
- ◆ The man pages for `sanlun` and the diagnostic scripts.

Supported Solaris, Data ONTAP features, configurations

Features supported by the Host Utilities

The Host Utilities support a number of features and configurations available with Solaris hosts and storage systems running Data ONTAP. Your specific environment affects what the Host Utilities support. Some of the supported features include:

- ◆ Multiple paths to the storage system when a multipathing solution is installed
- ◆ Veritas VxVM, Sun Volume Manager, and ZFS file system
- ◆ (MPxIO stack only) Asymmetrical Logical Unit Access
- ◆ Sun Microsystems' Logical Domains
- ◆ SAN booting

For information on which features are supported with which configurations, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

About the HBAs

The Host Utilities support a number of HBAs. The supported HBAs should be installed before you install the Host Utilities. Normally, they should have the correct firmware and FCode set. To determine the firmware and FCode setting on your system, run the appropriate administration tool for your HBA.

Note

For details on the specific HBAs that are supported and the required firmware and FCode values, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

If the Host Utilities setup procedure differs based on the HBA model, this guide describes the different steps. Unless otherwise noted, the information in this guide applies to all supported versions HBA models.

About multipathing

The Solaris Host Utilities support different multipathing solutions based on your configuration. Using multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, FC traffic continues on the remaining paths. In order for a host to have multiple paths to a LUN, you must have multipathing enabled.

The Veritas environment of the Host Utilities uses Veritas DMP to provide multipathing.

The MPxIO environment of the Host Utilities uses Sun's native multipathing solution.

About volume managers

The Solaris Host Utilities support different volume management solutions based on your stack.

The Veritas DMP stack uses Veritas Volume Manager (VxVM).

The MPxIO stack works with Sun Volume Manager (SVM) and VxVM to enable you to have different volume management solutions. In addition, you can use the ZFS file system to manage your volumes.

Note

To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.

(MPxIO only) ALUA support with certain versions of Data ONTAP

The MPxIO stack requires that you have Asymmetric Logical Unit Access (ALUA) enabled for active/active storage controllers. Stand-alone storage controllers provide symmetric access to LUNs and do not use ALUA.

Note

ALUA is also known as Target Port Group Support (TPGS).

ALUA defines a standard set of SCSI commands for discovering path priorities to LUNs on Fibre Channel and iSCSI SANs. When you have the host and storage controller configured to use ALUA, it automatically determines which target ports provide optimized (primary) and unoptimized (secondary) access to LUNs.

Check the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/> to see which versions of Data ONTAP support ALUA and are supported by the Host Utilities. NetApp introduced ALUA support with Data ONTAP 7.2 and Single Image cmode.

The supported cluster failover mode

If you upgrade to Data ONTAP 7.1 or later and have an active/active configuration (also known as a *storage cluster*), you should migrate the cmode setting of both systems in the configuration to single image mode. For

information on doing this, see *Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations*, which is available online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf.

About using Sun Microsystems Logical Domains

Certain configurations of the Host Utilities MPxIO stack support Sun Microsystems Logical Domains (LDom). The supported configurations include guests that are I/O Domains or guests that have iSCSI configured. You must install the Host Utilities if a guest is using NetApp storage.

If you are using LDom, you must configure your system with the NetApp settings. You can have the Host Utilities do this or you can configure the settings manually. For information on the NetApp settings, see either “[Veritas DMP: Configuring the Solaris Host](#)” on page 53 or “[MPxIO: Configuring the Solaris Host](#)” on page 65, depending on the configuration you are using.

A Solaris host running LDom accesses and uses NetApp storage exactly the same way any other Solaris host does.

For information on which configurations support LDom, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

About SAN booting

SAN boot is using a SAN-attached disk (a LUN), as a boot device for a Solaris host. Configuring SAN boot allows you to do the following:

- ◆ Use the SAN for your boot needs
- ◆ Use the SAN for both storage and boot needs, thereby consolidating and centralizing storage

The steps you need to perform to set up SAN booting differ based on your Host Utilities stack. If you have

- ◆ A Veritas DMP stack, see “[Veritas DMP: Configuring SAN Boot](#)” on page 159 for instructions on setting up SAN booting.
- ◆ A MPxIO stack, see “[MPxIO: Configuring SAN Boot](#)” on page 191 for instructions on setting up SAN booting.

Information on the stacks supported by the Host Utilities

Finding instructions for your Solaris environment

Many of the instructions in this manual apply to all the Solaris stacks that the Host Utilities support. In some cases, the commands you use vary based on a particular stack.

The following table summarizes areas where you must use instructions specific to your Solaris environment. If you are using a Veritas DMP stack, you must sometimes use different procedures based on whether you are using LPFC drivers or Sun native drivers.

Task	Veritas DMP configurations	MPxIO configurations
Executing commands	Log in as root and use standard Solaris and Veritas commands	Log in as root and use standard Solaris commands
Preparing the drivers	For information on setting the drivers up, see “ Veritas DMP: Managing the Driver Software ” on page 19	No manual steps are required. By default the drivers come with the operating system.
Configuring system parameters used by the Host Utilities	For information on the parameters used when you are running the Host Utilities with Veritas Storage Foundation, see “ Veritas DMP: Configuring the Solaris Host ” on page 53	For information on the parameters used when you are running the Host Utilities with the MPxIO stack, see “ MPxIO: Configuring the Solaris Host ” on page 65/
Special Setup for your environment	Install the ASL and APM. See “ Veritas DMP: Installing the ASL and APM ” on page 71.	See “ (MPxIO) Check the ALUA settings on the storage system ” on page 138
Working with LUNs	For information working with LUNs when using a Veritas DMP stack, see “ Veritas DMP: Configuring and Managing LUNs ” on page 99.	For information on working with LUNs when using a MPxIO stack, see “ MPxIO: Configuring and Managing LUNs ” on page 109.

Task	Veritas DMP configurations	MPxIO configurations
SAN booting	See “ Veritas DMP: Configuring SAN Boot ” on page 159 for instructions on setting up SAN booting.	For instructions on setting up SAN booting, see “ MPxIO: Configuring SAN Boot ” on page 191.

High-level look at Host Utilities Veritas DMP stack

The Host Utilities Veritas DMP stack works with Solaris hosts running Veritas Storage Foundation and DMP. The following is a summary of the supported stack at the time this document was produced.

Note

Check the NetApp Interoperability Matrix for details and current information about the supported stack.

- ◆ Operating system:
 - ❖ Solaris 9 and 10
- ◆ Processor:
 - ❖ SPARC processor systems
- ◆ HBA:
 - ❖ Emulex LPFC HBAs and their Sun-branded equivalents
 - ❖ Certain Sun OEM QLogic® HBAs and their Sun-branded equivalents
 - ❖ Certain Sun OEM Emulex® HBAs
- ◆ Drivers:
 - ❖ Emulex drivers (LPFC)
 - ❖ Sun-branded Emulex drivers (emlxs)
 - ❖ Sun-branded QLogic drivers (qlc)
- ◆ Multipathing
 - ❖ Veritas DMP

The Host Utilities Veritas DMP stack also supports the following:

- ◆ Volume manager
 - ❖ Veritas™ Volume Manager (VxVM)
- ◆ Clustering
 - ❖ Veritas Cluster Server (VCS)

High-level look at Host Utilities MPxIO stack

The Host Utilities MPxIO stack works with Solaris hosts running Sun StorEdge SAN Foundation Software and components that make up the native stack. The following is a summary of the supported stack at the time this document was produced.

Note

Check the NetApp Interoperability Matrix for details and current information about the supported stack.

- ◆ Operating system:
 - ❖ Solaris 10
- ◆ Processor:
 - ❖ SPARC processor systems
 - ❖ x86/64 (AMD) processor systems.
- ◆ HBA:
 - ❖ Certain QLogic HBAs and their Sun-branded equivalents
 - ❖ Certain Emulex HBAs and their Sun-branded equivalents
- ◆ Drivers:
 - ❖ Bundled Sun StorEdge SAN Foundation Software (SFS) drivers
 - ❖ Emulex drivers (emlxs)
 - ❖ Bundled Sun StorEdge SFS QLogic drivers (qlc)
- ◆ Multipathing
 - ❖ Sun StorageTek Traffic Manager (MPxIO)

The FC Solaris Host Utilities for MPxIO also support the following:

- ◆ Volume manager
 - ❖ SVM
 - ❖ ZFS
 - ❖ VxVM

Note

To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix, at <http://now.netapp.com/matrix/mtx/login.do/>.

- ◆ Clustering
 - ❖ Sun Clusters. This kit has been certified using the Sun Cluster Automated Test Environment (SCATE).
 - ❖ Veritas Cluster Server (VCS)

FC Solaris Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

Installing the FC Solaris Host Utilities and setting up your system involves a number of steps that are performed on both the storage system and the host. The following sections provide information about the tasks you need to perform.

Note

Occasionally there is a known problem that can affect your system setup. Please check the Solaris Host Utilities Release Notes to make sure there are no issues affecting your system setup. The Release Notes are updated whenever an issue is found and may contain more current information than this manual.

Prerequisites

Before you install the Host Utilities, you must perform the following tasks:

1. Set up your system:

- ❖ Host OS and appropriate updates
- ❖ HBAs
- ❖ Drivers
- ❖ **Veritas environment only:** Veritas Storage Foundation, the Array Support Library (ASL) for the storage controllers, and, with Veritas Storage Foundation 5.0, Array Policy Module (APM)

Note

You must install the Veritas VxVM before you install the ASL and APM software. These components are available from the Symantec web site.

- ❖ Volume management and multipathing, if used
- ❖ Storage system and Data ONTAP
- ❖ Switches (if used)
- ❖ Remove any version of the Host Utilities or Host Attach Kit that is currently installed

See the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/> for current information about the system requirements.

2. Verify that your storage system is:
 - ❖ Licensed correctly for the FC protocol service and running it
 - ❖ Using the recommended `cfmode`
 - ❖ **(MPxIO only)** Using ALUA
 - ❖ Set up to work with the correct volume manager for your stack
 - ❖ Configured to work with the target HBAs (if needed)
 - ❖ Set up to work with the Solaris host and the initiator HBAs
 - ❖ Set up with working volumes and qtrees (if desired)
3. If you're using a fabric connection, verify that the switch is:
 - ❖ Set up correctly
 - ❖ Zoned (if appropriate for your system)

Note

For information on supported topologies, see the *iSCSI and Fibre Channel Configuration Guide*, which is online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/fc_iscsi_config_guide.pdf.

- ❖ Cabled according to the instructions that came with the switch
4. Verify that:
 - ❖ The power to the switches, disk shelves, storage systems, and hosts is turned on in the correct order
 - ❖ The host and the storage system can communicate

Installation overview

Installing the Host Utilities involves the following tasks:

1. Get a copy of the Host Utilities software packages.
 - ❖ Get the compressed file containing the software package by either downloading it or copying it from the CD.
 - ❖ Extract the software packages for the supported stacks and the SAN toolkit from the compressed file.
2. Install the packages. You must install both the software package for your stack and the software package for the SAN toolkit:
 - ❖ Log in as root.
 - ❖ Go to the directory containing the extracted software package.

- ❖ Use the `pkgadd -d` command to install the Host Utilities package for your stack.
- ❖ Use the `basic_config` command to set the driver parameters.

Note

You can also set the parameters manually.

3. Configure the LUNs by performing the following tasks:

- ❖ Create at least one igroup and at least one LUN.
One way to create igroups and LUNs is to use the `lun setup` command.
- ❖ **(MPxIO only)** Enable ALUA, if you haven't already done so.
- ❖ Map the LUN to the igroup.
- ❖ **(LPFC drivers only)** Create persistent bindings using `hbacmd`. (You can also create persistent bindings using the Host Utilities-provided `create_binding.pl` script.)
- ❖ **(LPFC drivers 6.11cx2 and earlier only)** In the `sd.conf` file, create an entry for every NetApp LUN.
- ❖ Configure the host to discover the LUNs.
(LPFC drivers) Either perform a reconfiguration reboot (`touch /reconfigure; /sbin/init 6`) on the host or use the `/usr/sbin/devfsadm` command.
(Native drivers) Use the `/usr/sbin/cfgadm -c configure cx` command, where `x` is the controller number of the HBA where the LUN is expected to be visible.
- ❖ Label the LUNs using the Solaris format utility (`/usr/sbin/format`).
- ❖ Configure the volume management software
- ❖ Display information about the LUNs and HBA.
You can use the `sanlun` command.

For detailed information on these tasks

The following sections in this document provide more details about setting up your system, verifying your setup before you install the Host Utilities, installing and configuring the Host Utilities, and creating igroups and LUNs:

- ◆ **(LPFC drivers only)** [“Working with the Emulex LPFC software”](#) on page 24
- ◆ **(Native drivers only)** [“Working with the Sun Drivers for Emulex HBAs \(emlxs\)”](#) on page 34 and [“Working with the Sun Native Drivers for QLogic HBAs \(qlc\)”](#) on page 37

- ◆ [“Confirming that your system is set up correctly”](#) on page 128
- ◆ [“Installing the Host Utilities software package”](#) on page 42
- ◆ [“Veritas DMP: Configuring the Solaris Host”](#) on page 53
- ◆ [“MPxIO: Configuring the Solaris Host”](#) on page 65
- ◆ [“Installing or Upgrading the ASL and the APM”](#) on page 74
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 96
- ◆ [“Veritas DMP: Configuring and Managing LUNs”](#) on page 99
- ◆ [“MPxIO: Configuring and Managing LUNs”](#) on page 109
- ◆ [“Displaying host LUN information”](#) on page 116
- ◆ [“Displaying information about the storage system paths”](#) on page 119
- ◆ [“Displaying host HBA information”](#) on page 124

In addition, the section [“Where to go for more information”](#) on page 15 provides pointers to other documents and Web sites that can help you with your system setup.

Optional procedures

You may wish to perform the following tasks after installation is complete.

Task	Additional information
Configure SAN boot. SAN boot uses a SAN attached disk, like a LUN, as a boot device for a Solaris host	The instructions you use for setting up a SAN boot LUN vary based on your stack. See: ◆ “Veritas DMP: Configuring SAN Boot” on page 159 ◆ “MPxIO: Configuring SAN Boot” on page 191
Run diagnostic scripts. Normally, you would only run these scripts if you encountered a problem. These scripts collect information about your system and place it in compressed file that you can send to customer support.	See “Using Diagnostic Utilities” on page 140

Where to go for more information

Host and storage system information

The documents listed below contain information about host and storage system requirements, supported configurations, best practices, installation procedures and troubleshooting.

If you need information about...	See the following NetApp product documentation...
Known issues, troubleshooting, operational considerations, and post-release developments	The <i>Release Notes</i>, which are available at http://now.netapp.com/NOW/knowledge/docs/san/ Note— The <i>Release Notes</i> are updated more frequently than the rest of the documentation. You should always check the release notes before installing the Host Utilities to see if there have been any changes to the installation or setup process since this document was prepared. You should check them periodically to see if there is new information on using the Host Utilities. A summary of what has been updated and when is on the <i>Release Notes</i> index page.
The latest supported configurations	◆ The NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do/ ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
A summary of some of the commands you might use with the Host Utilities	The <i>Quick Reference</i> for the FC Solaris Host Utilities, which is available from the NOW site
Supported Fibre Channel SAN topologies and connecting the storage system to the fabric or switch	◆ Switch documentation at http://now.netapp.com/NOW/knowledge/docs/san/ ◆ <i>iSCSI and Fibre Channel Configuration Guide</i>, available on the NOW site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml
Changes to the host settings that are recommended by the Host Utilities	<i>Host Settings Affected by the Host Utilities</i> at http://now.netapp.com/NOW/knowledge/docs/san/

If you need information about...	See the following NetApp product documentation...
Changing the cfmode	<i>Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf
Configuring the storage system and managing SAN storage on it	◆ Index to Data ONTAP documentation at http://now.netapp.com/NOW/knowledge/docs/ontap ◆ <i>Best Practices for Reliability: New System Installation</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/index.shtml ◆ <i>Data ONTAP® Software Setup Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Block Access Management Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Release Notes</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>FilerView® Online Help</i> - available from within the FilerView GUI <i>Best Practices for Reliability: New System Installation</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/index.shtml
Enabling ALUA	<i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml In particular, see the section “<i>Enabling ALUA for a Fibre Channel igroup.</i>”
Verifying compatibility of a storage system with environmental requirements	<i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml

If you need information about...	See the following NetApp product documentation...
Upgrading Data ONTAP	<i>Data ONTAP® Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
Best practices/configuration issues	NetApp Knowledge Base at https://now.netapp.com/eservice/kbbrowse
Installing and configuring the HBA	Refer to your HBA vendor documentation.
Your host operating system and using its features, such as SVM, ZFS, or MPxIO	◆ Refer to your operating system documentation. You can download Sun manuals in PDF format from http://docs.sun.com/ .
Veritas Volume Manager	Refer to your Veritas documentation.
Working with Emulex	Refer to the Emulex documentation. The Emulex partner site for NetApp contains links to the documentation; see http://www.emulex.com/netapp/support/index.jsp
Working with QLogic	Refer to the QLogic documentation. QLogic provides a link to its NetApp partner site at http://support.qlogic.com/support/drivers_software.asp

Configuration overview

This chapter provides information on working with the drivers. It contains information about the following

- ◆ Emulex LPFC drivers
- ◆ Sun drivers for Emulex
- ◆ Sun drivers for QLogic

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Before you begin working with the drivers”](#) on page 20
- ◆ [“Working with the Emulex LPFC software”](#) on page 24
- ◆ [“Working with the Sun Drivers for Emulex HBAs \(emlxs\)”](#) on page 34
- ◆ [“Working with the Sun Native Drivers for QLogic HBAs \(qlc\)”](#) on page 37

Before you begin working with the drivers

Tasks to perform before you begin working with the drivers

Before you install new firmware or update FCodes, make sure you've taken care of the prerequisites:

- ◆ “[Verifying that system is set up correctly](#)” on page 20
- ◆ “[General information on getting the driver software](#)” on page 20
- ◆ Download the software for your HBA and extract the necessary files.
 - ❖ Emulex LPFC drivers and Sun-branded Emulex drivers: “[Downloading and extracting the Emulex software](#)” on page 21
 - ❖ Sun-branded QLogic drivers: “[Downloading and extracting the QLogic software](#)” on page 23

Verifying that system is set up correctly

Before you install the driver firmware, you must have set up:

- ◆ The storage system
 - ❖ Data ONTAP must be installed and running.
 - ❖ The storage system must be connected to the fabric or switch.
- ◆ The Solaris host
 - ❖ The host must be running the correct operating system and patches for the Host Utilities.
 - ❖ The HBA hardware must be installed.
 - ❖ The cables to the fabric or switch must be set up.

General information on getting the driver software

You can get the driver software from the company Web site for your HBA. To determine which drivers are supported with the Host Utilities, check the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

- ◆ Emulex LPFC HBAs: The Emulex software, including the Emulex utility programs and documentation, is available from the NetApp partner page on the Emulex site. Go to <http://www.emulex.com/ts/> and select NetApp. For information on downloading this software, see “[Downloading and extracting the Emulex software](#)” on page 21.
- ◆ Emulex HBAs with Sun native drivers: The Emulex software, including the Emulex utility programs and documentation, is available from the Solaris OS download section on the Emulex site. When this document was produced, that section was located at

<http://www.emulex.com/support/solaris/index.jsp>. For information on downloading this software, see “[Downloading and extracting the Emulex software](#)” on page 21.

- ◆ QLogic-branded HBAs: The QLogic SANsurfer CLI software and documentation are available on the QLogic support site. QLogic provides a link to its NetApp partner sites at http://support.qlogic.com/support/drivers_software.asp. You only need this software if you have to manipulate the FCode versions on QLogic-branded HBAs for SAN booting. For information on downloading this software, see “[Downloading and extracting the QLogic software](#)” on page 23.
- ◆ Sun-branded HBAs: You can also use certain Sun-branded HBAs. For more information on working with them, see the patch Readme file that Sun provides.

Downloading and extracting the Emulex software

Follow these steps to download and extract Emulex software and firmware.

Note

You may wish to download new firmware at the same time as you download the rest of the Emulex software. Your HBA may use an earlier version of the firmware than is recommended for the Host Utilities. To determine which version of the firmware you are currently running, see “[Determining the firmware version](#)” on page 35. To determine which firmware versions are supported, check the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Step	Action
1	On the Solaris host, create a download directory for the Emulex software. Enter: <code>mkdir /tmp/emulex</code>

Step	Action
2	To download Emulex driver and firmware, go to the location on the Emulex Web site for the type of drivers you are using: ◆ For Emulex HBAs using LPFC drivers, go to http://www.emulex.com/ts/ and click the NetApp link. Download the required driver, firmware, fcode, utilities, and user manuals. ◆ For Emulex HBAs using Sun native drivers, go to the Solaris OS download section. At the time this document was produced, that section was located at http://www.emulex.com/support/solaris/index.jsp Follow the instructions on the download page to get the driver software and place it in the /tmp/emulex directory. To determine which drivers are supported, see the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do.
3	On the Solaris host, change to the /tmp/emulex directory. <pre>cd /tmp/Emulex</pre>
4	Use the tar xvf command to extract the software from the files you downloaded. Note_____ If you are using Emulex Utilities for Sun native drivers, the .tar file you download contains two additional .tar files, each of which contains other .tar files. The file that contains the EMLXemlxu package for native drivers is emlxu_kit-<version>-sparc.tar. _____ Example: The following command lines show how to extract the software from the files. LPFC driver bundle: <pre>tar xvf solaris-<HBAnyware_version>-<driver_version>.tar</pre> Emulex Utility bundle for use with Sun Native drivers: <pre>tar xvf solaris-<HBAnyware_version>-<utility_version>-<subversion>.tar</pre>

Downloading and extracting the QLogic software

Follow these steps to download and extract QLogic software and firmware.

Step	Action
1	On the Solaris host, create a download directory for the QLogic software. Enter: mkdir /tmp/qlogic
2	To download the SANsurfer CLI software, go to http://www.qlogic.com and click the Downloads link.
3	Under “OEM Models,” click Netapp.
4	Click the link for your card type.
5	Choose the latest multiflash or bios image available and save it to the /tmp/qlogic directory on your host.
6	Change to the /tmp/qlogic directory and uncompress files that contain the SANsurfer CLI software package. Enter: uncompress scli-<version>.SPARC-X86.Solaris.pkg.Z

Working with the Emulex LPFC software

Information in this section

This section deals with the Emulex LPFC software. It provides information on the following tasks:

- ◆ [“Drivers prior to 6.21g: Installing the Emulex utilities for LPFC drivers”](#) on page 24
- ◆ [“Drivers prior to 6.21g: Unbinding the HBA”](#) on page 25
- ◆ [“Determining the Emulex LPFC firmware and FCode versions”](#) on page 32
- ◆ [“Upgrading the firmware and FCode”](#) on page 32

Emulex components

The Emulex software is provided as .tar files. You must download the correct .tar file for your system environment. The .tar files for Emulex HBAs with LPFC drivers contain the following components:

- ◆ The Host Bus Adapter driver software
- ◆ The HBA management suite, which includes HBAnyware and hbacmd
Sections of this document use hbacmd to install new firmware and to create persistent bindings. You can also perform these tasks using HBAnyware. If you prefer HBAnyware, refer to Emulex documentation for configuration information.
- ◆ Emulex Fibre Channel Utilities
For certain driver versions, the utility package is bundled with the driver. Beginning with the 6.21g version of the LPFC driver, Emulex started bundling a portion of the utilities with the LPFC driver itself. If you are using an LPFC driver version prior to 6.21g, the utility is a separately installed package bundle.

Drivers prior to 6.21g: Installing the Emulex utilities for LPFC drivers

After you extract the Emulex Utilities software bundle, follow the instructions in the Emulex *FCA Utilities Reference Manual* to install the package and then use the utilities.

Drivers prior to 6.21g: Adding the Emulex LPFC driver packages

After you extract the driver software bundle, follow the instructions in the Emulex user's manual that you downloaded to install the LPFC driver package only.

Do **not** install the HBAnyware packages at this time.

Do **not** reboot the host at this time, even if the manual instructs you to do so.

Drivers prior to 6.21g: Unbinding the HBA

Before you continue, unbind the Emulex HBA from the emlxs driver. Complete the following steps.

WARNING —
When you change the binding to the LPFC driver, the device paths for any LUNs attached to the HBA are changed, which could invalidate any /etc/vfstab entries that use those devices.

Step	Action
1	Unbind the HBA from the emlxs driver. Enter: <code>/opt/EMLXemlxu/bin/emlxdrv clear_emlxs</code> Result: The emlxs bindings are cleared.
2	Bind the HBA to the LPFC driver. Enter: <code>/opt/EMLXemlxu/bin/emlxdrv set_lpfc_nonsun</code> Result: The LPFC driver is bound to the HBAs.
3	Enter <code>q</code> to quit the emlxdrv tool.

Drivers prior to 6.21g: Removing the Emulex utilities

Because of a potential conflict between some versions of HBAnyware and the Emulex utilities, NetApp recommends that you uninstall the Emulex utilities before you continue. Follow the instructions in the Emulex *FCA Utilities Reference Manual* to remove the package.

Drivers prior to 6.21g: Adding the HBAnyware package

You should have extracted the driver and applications software bundle in `/tmp/emulex`. Follow the instructions in the Emulex user's manual that you downloaded to install the HBAnyware software package.

Drivers prior to 6.21g: Completing the driver installation

To finish the driver installation, you must shut down the host and ensure the adapters are in `set_sd_boot` mode. If the drivers are not in `set_sd_boot` mode, they will not bind to the LPFC driver when you boot up. Complete the following steps.

Step	Action
1	Set <code>auto-boot?</code> to false to prevent the host from booting up automatically. Enter: <code>eeeprom auto-boot?=false</code> Result: The PROM variable <code>auto-boot?</code> is set to false.
2	Halt the system. Enter: <code>/sbin/init 0</code> Result: The host is taken down to the <code>OK></code> prompt.
3	Reset the host. Enter: <code>OK> reset_all</code> Result: The host resets and runs the power-on self test. It stops at the <code>OK></code> prompt.

Step	Action
4	Run the show-devs command to find the device paths for the HBAs. Note When you run this command, be sure you document the paths for your HBAs. Result: When you run this command, a list of the HBA paths appears. Be sure you document the paths for your HBAs. These paths could show up as lpfc@, emlxs@, or fibre-channel@. If they show up as emlxs@, continue with step 5; otherwise you can skip to step 8. Example: The show-devs command produces output similar to the following. <pre>ok> show-devs controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/lpfc@3 /pci@8,700000/lpfc@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3</pre>
5	Switch the HBA to set-sd-boot mode. Enter: <pre>OK> select <device path> OK> set-sd-boot</pre>
6	Repeat step 5 for each HBA.

Step	Action
7	Reset the host so the boot mode changes are implemented. Enter: OK> reset_all Result: The host rests and runs the power on self test. It stops at the OK> prompt.
8	Enable auto-boot mode. Enter: OK> setenv auto-boot? true Result: Auto Boot mode is turned back on.
9	Boot the host. Enter: OK> boot -r Result: The host performs a reconfiguration boot.

Drivers 6.21g and later: Adding the Emulex LPFC and HBAnyware driver packages

After you extract the driver software bundle, follow the instructions in the Emulex user's manual that you downloaded to install the LPFC driver and HBAnyware packages.

Do **not** reboot the host at this time, even if the Emulex manual instructs you to do so.

Drivers 6.21g and later: Unbinding the HBA

Before you continue, unbind the Emulex HBA from the emlxs driver. Complete the following steps.

WARNING

When you change the binding to the LPFC driver, the device paths for any LUNs attached to the HBA are changed, which could invalidate any /etc/vfstab entries that use those devices.

Step	Action
1	Unbind the HBA from the emlxs driver. Enter: /opt/EMLXemlxu/bin/emlxdrv clear_emlxs Result: The emlxs bindings are cleared.

Step	Action
2	Bind the HBA to the LPFC driver. Enter: <code>/opt/EMLXemlxu/bin/emlxdrv set_lpfc_nonsun</code> Result: The LPFC driver is bound to the HBAs.
3	Enter <code>q</code> to quit the emlxdrv tool.

Drivers 6.21g and later: Completing the driver installation

To finish the driver installation, you must shut down the host and ensure the adapters are in `set_sd_boot` mode. Complete the following steps.

Step	Action
1	Set <code>auto-boot?</code> to false to prevent the host from booting up automatically. Enter: <code>eeeprom auto-boot?=false</code> Result: The PROM variable <code>auto-boot?</code> is set to false.
2	Halt the system. Enter: <code>/sbin/init 0</code> Result: The host is taken down to the <code>OK></code> prompt.
3	Reset the host. Enter: <code>OK> reset_all</code> Result: The host rests and runs the power-on self test. It stops at the <code>OK></code> prompt.

Step	Action
4	Run the show-devs command to find the device paths for the HBAs. Note When you run this command, be sure you document the paths for your HBAs. These paths could show up as "lpfc@", "emlxs@", or "fibre-channel@". If they show up as "emlxs@", continue with step 5; otherwise you can skip to step 8. Result: When you run this command, a list of the HBA paths appears. Be sure you document the paths for your HBAs. These paths could show up as lpfc@, emlxs@, or fibre-channel@. If they show up as emlxs@, continue with step 5; otherwise you can skip to step 8. Example: The show-devs command produces output similar to the following. <pre>ok> show-devs controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/lpfc@3 /pci@8,700000/lpfc@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3</pre>
5	Switch the HBA to set-sd-boot mode. Enter: <pre>OK> select <device path> OK> set-sd-boot</pre>
6	Repeat step 5 for each HBA.

Step	Action
7	Reset the host so the boot mode changes are implemented. Enter: OK> reset_all Result: The host rests and runs the power-on self test. It stops at the OK> prompt.
8	Enable auto-boot mode. Enter: OK> setenv auto-boot? true Result: Auto Boot mode is turned back on.
9	Boot the host. Enter: OK> boot -r Result: The host performs a reconfiguration boot.

Special Note about the Emulex LPFC driver utility hbacmd

The LPFC driver utility hbacmd is located in two different places. The version you download depends on your driver version:

- ◆ For lpfc driver version 6.21g and later, the full path is /opt/HBAnyware/hbacmd.
- ◆ For lpfc driver versions prior to 6.21g, the full path is /usr/sbin/hbanyware/hbacmd.

This document does not use the full path in the examples. Please use the appropriate path based on the driver version you are installing.

Determining the Emulex LPFC firmware and FCode versions

If you are not using the Emulex firmware recommended for the Host Utilities, you must upgrade your firmware. To determine which version of firmware you should be using and which version you are actually using, complete the following step.

Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current firmware requirements. Once you know the requirements, you can use <code>hbacmd</code> to determine if your firmware is current or needs updating.
2	Use <code>hbacmd</code> to document the current installed adapters. Enter: <code>hbacmd listhbas</code> Result: The software displays the list of installed adapters and their World Wide Port Names (WWPN).
3	Use <code>hbacmd</code> to determine the firmware of each adapter. Enter: <code>hbacmd hbaattributes <adapter_wwpn></code> Result: The software displays information about the adapter.

Upgrading the firmware and FCode

To upgrade your firmware version, complete the following steps:

Note

You can download the latest firmware from the NetApp page on www.emulex.com/ts/index.htm.

Step	Action
1	Use <code>hbacmd</code> to document the current installed adapters. Enter: <code>hbacmd listhbas</code> Result: The software displays the list of installed adapters and their WWPNs.

Step	Action
2	Use hbacmd to update the firmware. You must use this command with each adapter. Enter: hbacmd download <adapter_wwpn> <firmware_filename> Result: The firmware is upgraded on the adapter.
3	Use hbacmd to update the fcode. This command should be repeated for each adapter. Enter: hbacmd download <adapter_wwpn> <fcode_filename> Result: FCode is upgraded on the adapter.

Working with the Sun Drivers for Emulex HBAs (emlxs)

Information in this section

This section deals with the Emulex-branded HBAs and Sun-branded Emulex OEM HBAs that use the native driver (emlxs) software. It provides information on the following tasks:

- ◆ “[Emulex components](#)” on page 34
- ◆ “[Installing the EMLXemlxu utilities](#)” on page 34
- ◆ “[Determining the firmware version](#)” on page 35
- ◆ “[Upgrading the firmware](#)” on page 35
- ◆ “[Upgrading the firmware](#)” on page 35

Emulex components

The Emulex software is provided as .tar files. You must download the correct .tar file for your system environment. When you are using Emulex HBAs with native drivers, you need the .tar files containing the Emulex Fibre Channel Adapter (FCA) Utilities (EMLXemlxu).

You use these utilities to manage the firmware and FCode of the Emulex HBAs with native driver. To install and use these utilities, follow the instructions in the Emulex *FCA Utilities Reference Manual*.

Installing the EMLXemlxu utilities

After you extract the EMLXemlxu utilities, run the emlxs_install command to install the EMLXemlxu package:

```
# ./emlxs_install
```

Note

For more information on installing and using these utilities, see the Emulex *FCA Utilities Reference Manual*.

Determining the firmware version

If you are not using the Emulex firmware recommended for the Host Utilities, you must upgrade your firmware. To determine which version of firmware you should be using and which version you are actually using, complete the following steps:

Note-Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current firmware requirements. Once you know the requirements, you can use <code>emlxadm</code> to determine if your firmware is current or needs updating.
2	Run the <code>emlxadm</code> utility. Enter: <code>/opt/EMLXemlxu/bin/emlxadm</code> Result: The software displays a list of available adapters.
3	Select the device you want to check. Result: The software displays a menu of options.
4	Exit the <code>emlxadm</code> utility by entering <code>q</code> at the <code>emlxadm></code> prompt.

Upgrading the firmware

Follow the steps below to upgrade the firmware for your HBA.

Note _____
Sun-branded HBAs have the proper firmware version pushed to the card by the native driver.

Step	Action
1	Run the <code>emlxadm</code> utility. Enter: <code>/opt/EMLXemlxu/bin/emlxadm</code> Result: The software displays a list of available adapters.
2	Select the device you want to check. Result: The software displays a menu of options.

Step	Action
3	At the emlxadm> prompt, enter: download_fw <filename> Result: The firmware is loaded onto the selected adapter.
4	Exit the emlxadm utility by entering q at the emlxadm> prompt.
5	Reboot your host.

Updating your FCode

Follow the steps below to upgrade the FCode for your HBA.

Step	Action
1	Run the emlxadm utility. Enter: /opt/EMLXemlxu/bin/emlxadm Result: The software displays a list of available adapters.
2	Select the device you want to check. Result: The software displays a menu of options.
3	At the emlxadm> prompt, enter: download_fcode <filename> Result: The FCode is loaded onto the selected adapter.
4	Exit the emlxadm utility by entering q at the emlxadm> prompt.

Working with the Sun Native Drivers for QLogic HBAs (qlc)

Information in this section

This section deals with the QLogic-branded and Sun-branded QLogic OEM HBAs that use the native driver (qlc) software. It provides information on the following tasks:

- ◆ [“Installing the SANsurfer CLI package”](#) on page 37
- ◆ [“Determining the FCode on QLogic cards”](#) on page 37
- ◆ [“Upgrading the QLogic FCode”](#) on page 38

Installing the SANsurfer CLI package

After you extract the QLogic files, you need to install the SANsurfer CLI package. To do this, complete the following steps.

Step	Action
1	Install the SANsurfer CLI package using the pkgadd command. Enter: pkgadd -d /tmp/qlogic/scli-<version>.SPARC-X86.Solaris.pkg
2	From the directory where you extracted the QLogic software, unzip the FCode package. Enter: unzip <fcode_filename>.zip
3	For instructions on updating the FCode, please see the “Upgrading the QLogic FCode” on page 38.

Determining the FCode on QLogic cards

If you are not using the FCode recommended for the Host Utilities, you must upgrade it. To determine the FCode you should be using and which one you are actually using, complete the following steps:

Step	Action
1	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the current FCode requirements. Now, use scli utility to determine if your firmware is current or needs updating.

Step	Action
2	Run the scli utility. Enter: <code>/usr/sbin/scli</code> Result: The software displays a menu.
3	Select option 3 (HBA Information Menu). Result: The software displays the HBA Information Menu.
4	Select option 1 (Information). Result: The software displays a list of available ports.
5	Select the adapter port for which you want information. Result: The software displays information about that HBA port.
6	Write down the FCode version and press Return. Result: The software displays a list of available ports.
7	Repeat steps 5 and 6 for each adapter you want to query. When you have finished, select option 0 to return to the main menu. Result: The software displays the main menu.
8	To exit the scli utility, select option 13 (Quit).

Upgrading the QLogic FCode

To upgrade the QLogic FCode, complete the following steps.

Step	Action
1	Run the scli utility. Enter: <code>/usr/sbin/scli</code> Result: The software displays a menu.
2	Select option 8 (HBA Utilities). Result: The software displays a menu.

Step	Action
3	Select option 3 (Save Flash). Result: The software displays a list of available adapters.
4	Select the number of the adapter for which you want information. Result: The software displays a file name to use.
5	Enter the name of the file into which you want to save the flash contents. Result: The software backs up the flash contents and then waits for you to press Return.
6	Press Return. Result: The software displays a list of available adapters.
7	If you are upgrading more than one adapter, repeat steps 4 through 6 for each adapter. When you are done, select option 0 to return to the main menu.
8	Select option 0 to return to the main menu.
9	Select option 8 (HBA Utilities). Result: The software displays a menu.
10	Select option 1 (Update Flash). Result: The software displays a menu of update options.
11	Select option 1 (Select an HBA Port) Result: The software displays a list of available adapters
12	Select the appropriate adapter number. Result: The software displays a list of Update ROM options
13	Select option 1 (Update Option ROM). Result: The software requests a file name to use.

Step	Action
14	Enter the file name of the multiflash firmware bundle that you extracted from the file you downloaded from QLogic. The file name should be similar to q24mf129.bin Result: The software upgrades the FCode.
15	Press Return. Result: The software displays a menu of update options.
16	If you are upgrading more than one adapter, repeat steps 11 through 15 for each adapter.
17	When you have finished, select option 0 to return to the main menu.
18	To exit the scli utility, select option 13 (Quit).

Chapter overview This chapter explains how to install the FC Solaris Host Utilities.

Topics in this chapter This chapter includes the following topics:

- ◆ “[Installing the Host Utilities software package](#)” on page 42
- ◆ “[Upgrading the FC Solaris Host Utilities](#)” on page 48
- ◆ “[Removing an FC Solaris Host Utilities or Attach Kit installation](#)” on page 49

Installing the Host Utilities software package

Before you begin

Before you install the Host Utilities SAN Toolkit software package:

- ◆ Verify that your host system meets requirements and is set up correctly. To determine the current hardware and software requirements for the Host Utilities, search the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/> for the latest hardware and software requirements.
- ◆ If you currently have the FC Solaris Host Utilities or FC Solaris Attach Kit installed, you must remove it. See “[Removing an FC Solaris Host Utilities or Attach Kit installation](#)” on page 49.

Key steps involved in setting up the Host Utilities

Once your host is set up, you must install the Host Utilities software package for your stack.

When you have installed the software, you can use the tools it provides to complete your setup and configure your host parameters. The Host Utilities provide the following configuration tools:

- ◆ **(MPxIO only)** `mpxio_set` is used if you have used the iSCSI Solaris Host Utilities 3.0.1 or earlier. This script removes the symmetric options from the `scsi_vhci.conf` file. See “[Enable ALUA for MPxIO](#)” on page 66.
- ◆ `basic_config` configures the host system parameters. For more information, see either the Veritas DMP information on using `basic_config` in “[Overview of using the basic_config command](#)” on page 56 or the MPxIO information in “[Using basic_config to configure system parameters for MPxIO stacks](#)” on page 67.
- ◆ **(Veritas DMP only)** `create_binding.pl` creates persistent bindings for systems using LPFC drivers. For more information, see “[Veritas DMP: Using the create_binding.pl Script](#)” on page 149.
- ◆ `sanlun` displays host and storage system information. For more information, see “[Using sanlun](#)” on page 115.

(Veritas DMP only) You can also use tools from the HBA Management Suite to perform these tasks. If you want to use HBA management suite tools, you do not need to install the FC Solaris Host Utilities.

After you install the Host Utilities software, you will need to configure the host system parameters:

- ◆ For Veritas DMP stacks, see “[Veritas DMP: Configuring the Solaris Host](#)” on page 53
- ◆ For MPxIO stacks, see “[MPxIO: Configuring the Solaris Host](#)” on page 65

Obtaining the software

The software is provided as compressed files on a CD-ROM. These files are also available for downloading from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software>. There are two Host Utilities software distribution packages:

- ◆ SPARC processor systems. This is the software package you must install if you have either a Veritas DMP stack or a MPxIO stack using a SPARC processor.
- ◆ x86/64 (AMD) systems. This is the package you install if you have a MPxIO stack that is using an x86/64 (AMD) processor.

You only need to install the file that is appropriate for your system.

Obtaining the compressed file from the CD-ROM

To obtain the compressed file containing your software package from the CD-ROM, complete the following steps.

Step	Action
1	Insert the CD-ROM into the CD-ROM drive.
2	Change to the root directory of the CD-ROM: <code># cd <CDROM_mountpoint></code>
3	Copy the compressed file for your stack from the CD-ROM's root directory to a temporary location on the host.

Downloading the software from the NOW site

To download the software from the NOW site, complete the following steps.

Step	Action
1	Log in to the NOW site at http://now.netapp.com and select the Download Software page.

Step	Action
2	On the Download Software page, go to the FC Host Utilities (Attach Kits) product row of the table and select Solaris from the Select Platform drop-down list and click Go.
3	Click Go.
4	Follow the prompts to reach the Software Download page.
5	Download the compressed file to a local directory on your host. Note —— If you download the file to a machine other than the Solaris host, you must move it to the host.

Installing the Host Utilities software

To install the Host Utilities software, complete the following steps.

Note——

It is a good practice to check the Solaris Host Utilities Release Notes to see if there have been any changes or new recommendations for installing and using the Host Utilities since this installation guide was produced.

Step	Action
1	Login to the host system as root.

Step	Action
2	Place the compressed file for your processor in a directory on your host and go to that directory. At the time this documentation was prepared, the compressed files were called: ◆ SPARC CPU: <code>netapp_solaris_host_utilities_5_0_sparc.tar.gz</code> ◆ x86/x64 (AMD) CPU: <code>netapp_solaris_host_utilities_5_0_amd.tar.gz</code> Note_____ The actual file names for the Host Utilities software may be slightly different from the ones shown in these steps. _____ Example: If you are installing the <code>netapp_solaris_host_utilities_5_0_sparc.tar.gz</code> file on a SPARC system, you might put it in the <code>/tmp</code> directory on your Solaris host. You can use the following commands to place the file in that directory and then move to the directory: <pre># cp netapp_solaris_host_utilities_5_0_sparc.tar.gz /tmp # cd /tmp</pre>
3	Unzip the file using the <code>gunzip</code> command. Result: The software unzips the <code>tar.gz</code> files Example: The following example uses <code>gunzip</code> to unzip the files for a SPARC system. <pre># gunzip netapp_solaris_host_utilities_5_0_sparc.tar.gz</pre>
4	Untar the file. You can use the <code>tar xvf</code> command to do this. Result: The Host Utilities scripts are extracted to the default directory. Example: The following example uses <code>tar xvf</code> to extract the Solaris installation package for a SPARC system. <pre># tar xvf netapp_solaris_host_utilities_5_0_sparc.tar</pre>
5	Add the packages that you extracted from tar file to your host. You can use the <code>pkgadd</code> command to do this. Result: The packages are added to the <code>/opt/NTAP/SANToolkit/bin</code> directory. Example: The following example uses the <code>pkadd</code> command to install the Solaris installation package. <pre># pkgadd -d NTAPSANTool.pkg</pre>

Step	Action
6	Confirm that the toolkit was successfully installed by using the pkginfo command. Example: This example uses the pkginfo command to confirm that the toolkit was successfully installed. <pre># pkginfo NTAPSANTool application NTAPNTAPSANTool Solaris SAN Tool Utility for NetApp filers.</pre> Once the toolkit is installed, you should have the following tools in the installation directory: <pre>[20] 11:00:50 (root@shu05) /opt/NTAP/SANToolkit/bin # ls -la total 3458 drwxrwxr-x 2 root sys 512 Nov 13 11:00 ./ drwxrwxr-x 3 root sys 512 Nov 13 11:00 ../ -r-xr-xr-x 1 root sys 21146 Nov 10 16:28 basic_config* -r-xr-xr-x 1 root sys 4733 Nov 10 16:28 brocade_info* -r-xr-xr-x 1 root sys 5229 Nov 10 16:28 cisco_info* -r-xr-xr-x 1 root sys 52670 Nov 10 16:28 create_binding.pl* -r-xr-xr-x 1 root sys 9543 Nov 10 16:28 filer_info* -r-xr-xr-x 1 root sys 43123 Nov 10 16:28 FTP.pm* -r-xr-xr-x 1 root sys 5287 Nov 10 16:28 mcdata_info* -r-xr-xr-x 1 root sys 14052 Nov 10 16:27 mpxio_set* -r-xr-xr-x 1 root sys 6813 Nov 10 16:28 qllogic_info* -r-xr-xr-x 1 root sys 996 Nov 10 16:28 san_version* -r-xr-xr-x 1 root sys 1394708 Nov 10 16:28 sanlun* -r-xr-xr-x 1 root sys 16032 Nov 10 16:28 SHsupport.pm* -r-xr-xr-x 1 root sys 20169 Nov 10 16:28 solaris_info* -r-xr-xr-x 1 root sys 12772 Nov 10 16:28 ssd_config.pl* -r-xr-xr-x 1 root sys 130743 Nov 10 16:28 Telnet.pm* -r-xr-xr-x 1 root sys 677 Nov 10 16:28 vidpid.dat*</pre> Note_____ For information on these files, see “Software tools provided by the Host Utilities” on page 3.
7	To complete the installation, you must configure the host parameters. For Veritas DMP stacks, see “Veritas DMP: Configuring the Solaris Host” on page 53 For MPxIO stacks, see “MPxIO: Configuring the Solaris Host” on page 65

**Obtaining
documentation for
the Host Utilities**

Documentation for the Host Utilities is available on the NOW site and on the FC Solaris Host Utilities 5.0 CD-ROM. The SAN/iSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>) contains documentation for all the NetApp Host Utilities.

Upgrading the FC Solaris Host Utilities

Upgrading the Host Utilities or reverting to another version

You can upgrade to a newer version of the Host Utilities or revert to a 4.x or 3.0 version of the Host Utilities without any effect on system I/O.

To upgrade or revert, complete the following steps:

Step	Action
1	Use the Solaris <code>pkgrm</code> command to remove the Host Utilities software package you no longer need. For more information on removing the Host Utilities, see “Removing an FC Solaris Host Utilities or Attach Kit installation” on page 49. Note _____ Removing the software package does not remove or change the system parameter settings for that I/O stack. To remove the settings you added when you configured the Host Utilities, you must perform additional steps. You do not need to remove the settings if you are upgrading the Host Utilities. _____
2	Use the Solaris <code>pkgadd</code> command to add the appropriate Host Utilities software package. For steps on installing and configuring the Host Utilities, see “Installing the Host Utilities software package” on page 42 and “” on page 51.

Removing an FC Solaris Host Utilities or Attach Kit installation

Methods for removing the Host Utilities or Attach Kit

You should never install the Host Utilities on a host that currently has a version of the FC Solaris Host Utilities or Attach Kit installed. You must remove the old version before installing a new version or reverting to a different version.

The method you use for uninstalling the Host Utilities or Attach Kit depends on the version of the kit that you installed.

- ◆ For FC Solaris Host Utilities 4.x, 3.0, use the `pkgrm` command to remove the Host Utilities software package.
- ◆ For FC Solaris Attach Kit 2.0, use the `uninstall` script from the FC Solaris Attach Kit 2.0 software to uninstall the Attach Kit software.

Removing FC Solaris Host Utilities 4.x, 3.0

To remove the FC Solaris Host Utilities 4.x, 3.0, complete the following steps.

Step	Action
1	Remove the FC parameters that were set when you ran the <code>basic_config</code> command or that you set manually. To restore the values, you can do one of the following: ◆ Replace the system files with the backups you made before changing the values ❖ SPARC systems and systems using Sun native drivers: <code>/kernel/drv/ssd.conf</code>. ❖ (MPxIO) x86/64 (AMD) systems: <code>/kernel/drv/sd.conf</code>. ❖ (Veritas DMP) LPFC drivers: replace <code>/etc/system</code>, <code>/kernel/drv/lpfc.conf</code>, and <code>/kernel/drv/sd.conf</code>. ◆ Use the <code>basic_config</code> command to revert to the saved values (you can only do this once). ❖ SPARC systems and systems using Sun native drivers: <code>basic_config -ssd_unset</code> ❖ (MPxIO) x86/x64 (AMD) systems: <code>basic_config -sd_unset</code> ❖ (Veritas DMP) LPFC drivers: use the <code>basic_config -r</code> command.

Step	Action
2	Use the <code>pkgrm</code> command to remove the FC Solaris Host Utilities software from the <code>/opt/NTAP/SANToolkit/bin</code> directory. Example: The following command line removes the Host Utilities software package. <pre># pkgrm NTAPSANTool</pre>
3	(MPxIO only) You can disable MPxIO by using the <code>stmsboot</code> command. This command is only available with the following systems: ◆ SPARC ◆ x86/x64 (AMD) running Solaris 10 Update 4 Example: The following command line disables MPxIO. <pre># stmsboot -d</pre>
4	To enable the changes, reboot your system. Example: You can use the following command to reboot your system. <pre># touch /reconfigure # init 6</pre>

Uninstalling the Attach Kit 2.0 software

To uninstall the Attach Kit 2.0 software, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Locate the FC Solaris Attach Kit 2.0 software. By default, the FC Solaris Attach Kit is installed in <code>/opt/NTAPsanlun/bin</code> .

Step	Action
3	From the /opt/NTAPsanlun/bin directory, enter the <code>./uninstall</code> command to remove the existing software. Example: You can use the following command to uninstall the existing software. <pre># ./uninstall</pre> Note _____ The uninstall script automatically creates a backup copy of the /kernel/drv/lpfc.conf and sd.conf files as part of the uninstall procedure. It is a good practice, though, to create a separate backup copy before you begin the uninstall.
4	At the prompt “Are you sure you want to uninstall lpfc and sanlun packages?” enter y. Result: The uninstall script creates a backup copy of the /kernel/drv/lpfc.conf and sd.conf files to /usr/tmp. The backup files are named as follows: ◆ lpfc.conf.save ◆ sd.conf.save If a backup copy already exists, the install script prompts you to overwrite the backup copy.
5	Reboot your system. Example: You can use the following command to reboot your system. <pre># touch /reconfigure # init 6</pre>

Chapter overview

This chapter explains how to configure the Solaris host for the Veritas DMP stack. It provides information on the system parameters that need to be changed and methods for changing them.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Completing the installation and configuring system parameters](#)” on page 54
- ◆ “[Overview of using the basic_config command](#)” on page 56
- ◆ “[Configuring parameters on systems using LPFC drivers](#)” on page 57
- ◆ “[Configuring parameters on systems using Sun native drivers](#)” on page 62

Obtaining documentation

Documentation for the Host Utilities is available on the NOW site and on the FC Solaris Host Utilities 5.0 CD-ROM. The SAN/IPSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>) contains documentation for all the NetApp Host Utilities.

Completing the installation and configuring system parameters

Completing the Host Utilities setup on a Veritas DMP stack

To complete the installation when you're using a Veritas DMP stack, you must configure the system parameters. The tasks you perform vary slightly depending on your driver.

- ◆ For LPFC drivers, see [“Configuring parameters on systems using LPFC drivers”](#) on page 57.
To configure these parameters, you must modify the `/etc/system` file and `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf` files.
- ◆ For Sun native drivers, see [“Configuring parameters on systems using Sun native drivers”](#) on page 62.
To configure these parameters, you must modify the `/kernel/drv/ssd.conf` file.

There are two ways to modify these files:

- ◆ Manually editing them.
- ◆ Using the `basic_config` command to modify them. This command is provided as part of the FC Solaris Host Utilities and automatically sets these files to the correct values.

Note

The `basic_config` command does not modify the `/kernel/drv/sd.conf` file unless you are using an x86/x64 (AMD) processor with MPxIO. For information on that configuration, see [“MPxIO: Configuring the Solaris Host”](#) on page 65.

Note

For more information on the `basic_config` command, see [“Overview of using the basic_config command”](#) on page 56. For a complete list of the host parameters that NetApp recommends you change and an explanation of why those changes are recommended, see the *Host Settings Affected by the Host Utilities* document, which is on the NOW site on the SAN/IPSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>).

Before you begin configuring the Host Utilities

Complete the following prerequisites before you configure system parameters.

- ◆ (LPFC drivers only) Enter the `fcpl show cfmodes` command on the storage system to identify the storage system's cfmodes setting. Required values in the host configuration files vary, depending on the cfmodes setting.

If you plan to change the storage system cfmodes to single image, first complete the procedure in *Changing the Cluster cfmodes Setting in Fibre Channel SAN Configurations*, which is available online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf. Then follow the procedure below to configure system parameters.

- ◆ Create your own backup of the files you are modifying:

For systems using LPFC drivers, make backups of these files:

- ❖ `/etc/system`
- ❖ `/kernel/drv/lpfc.conf`
- ❖ `/kernel/drv/sd.conf`

For systems using Sun native drivers, make a backup for this file:

- ❖ `/kernel/drv/ssd.conf`

The `basic_config` command automatically creates backups for you, but you can only revert to those backups once. By manually creating the backups, you can revert to them as needed.

Overview of using the `basic_config` command

Using the `basic_config` command

The `basic_config` command is a configuration script that you can use to automatically set system values to the ones recommended for use with the Host Utilities.

You execute the `basic_config` command from the command line. The options you supply with this command depend on whether your system environment uses LPFC or Sun native drivers.

Tasks performed by the `basic_config` command

You can use the `basic_config` command to set system values to the ones recommended for use with the Host Utilities.

- ◆ For systems using LPFC drivers, the `basic_config` command can make the following modifications:
 - ❖ Update items in `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf` to the values recommended for NetApp storage systems using LPFC drivers.
 - ❖ Update time-out parameters in the `/etc/system` and `/kernel/drv/lpfc.conf` file, based on the `cfmode` that you specify.
 - ❖ Identify your Solaris OS version and set recommended values for that version in the `/etc/system` file.

See [“Configuring parameters on systems using LPFC drivers”](#) on page 57.

- ◆ For systems using Sun native drivers, the `basic_config` command can update items in `/kernel/drv/ssd.conf` to the values recommended for NetApp storage systems using Sun native drivers. See [“Configuring parameters on systems using Sun native drivers”](#) on page 62.

On systems using LPFC drivers, the `basic_config` command displays the recommended values for your approval before the file is updated. It also creates backup copies of the `/etc/system` and `/kernel/drv/lpfc.conf` or `/kernel/drv/ssd.conf` files.

Configuring parameters on systems using LPFC drivers

Tasks involved in configuring systems using LPFC drivers

If your system uses LPFC drivers, you must complete the configuration by performing the following tasks:

- ◆ Identify the cfmode setting of the storage system
- ◆ Set /etc/system parameters to required values
- ◆ Set /kernel/drv/lpfc.conf to the required values

You can manually set the values in these files or you can set them automatically by running the basic_config command. The sections that follow describe these tasks.

Note

You must also set up persistent bindings and create an entry in the sd.conf file. For information on setting up persistent bindings, see “[Veritas DMP: Using the create_binding.pl Script](#)” on page 149.

Identifying the cfmode

The values you supply for certain parameters vary depending on the storage system’s cfmode. To identify the cfmode, complete the following steps.

Step	Action
1	Access the storage system. You can do this by using a command such as telnet or going directly to the storage system.
2	On the storage system enter the following command: fcpl show cfmode Result: The storage system displays its cfmode.

Setting lpfc.conf variables manually on systems using LPFC drivers

You can manually set the /kernel/drv/lpfc.conf variables so that they match the storage system cfmode. To do this, first determine the storage system’s cfmode (see “[Identifying the cfmode](#)” on page 57). Now update the file by entering the appropriate values:

- ◆ For Standby (clustered storage systems)
 - ❖ linkdown-tmo=30

- ❖ nodev-tmo=**120**
- ◆ For Single_image or Partner (clustered storage systems)
 - ❖ linkdown-tmo=**30**
 - ❖ nodev-tmo=**30**
- ◆ For Stand Alone (single storage system)
 - ❖ linkdown-tmo=**255**
 - ❖ nodev-tmo=**255**

The default value for the topology parameter is 0. You can reset it to either of the following:

- ◆ topology=**2** - for switch (point to point)
- ◆ topology=**4** - for arbitrated loop

To set values for all other parameters, enter:

- ◆ tgt-queue-depth=**256**
- ◆ lun-queue-depth=**16**
- ◆ scan-down=**0**
- ◆ num-iocbs=**1024**
- ◆ num-bufs=**1024**
- ◆ dqfull-throttle-up-time=**4**
- ◆ dqfull-throttle-up-inc=**2**

The Solaris Host Utilities provides a best-fit setting for target and/or LUN queue depths. NetApp provides documentation for determining host queue depths and sizing. For information on NetApp FC and iSCSI storage systems, see:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

For information on configuration limitations associated with NetApp storage controllers, see the *iSCSI and Fibre Channel Configuration Guide*:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/config_guide/setting_up/concept/c_oc_set_config-limits-overview.html

Setting /etc/system variables manually on systems using LPFC

You can manually set the /etc/system variables so that they match the storage system cfmode. To do this, first determine the storage system's cfmode (see “[Identifying the cfmode](#)” on page 57). Now update the file by entering one of the following:

- ◆ For Standby (clustered storage systems), set the value to 120 seconds:
 - ❖ set sd:sd_io_time=**0x78**

- ◆ For Single_image or Partner (clustered storage systems), set the value to 30 seconds:
 - ❖ set sd:sd_io_time=0x1e
- ◆ For Stand-Alone (single storage systems), set the value to 255 seconds:
 - ❖ set sd:sd_io_time=0xff

If your host is Solaris 9 and you are configuring the FC Solaris Host Utilities, set the following:

- ◆ set sd:sd_retry_count=5

Using basic_config to set the options automatically on systems using LPFC drivers

On systems using LPFC drivers, the basic_config command has the several options you can use. The following are the acceptable basic_config command lines:

```
basic_config -h
basic_config -p
basic_config -r
basic_config -v
basic_config -m <cfmode> [-l|-s]
```

The following table explains these options. For an example of using the basic_config command, see [“Example of using basic_config on systems with LPFC drivers”](#) on page 60.

Option	Description	Notes
-h, -help	Displays the usage page.	~
-l	For direct attached configurations, sets the topology parameter in the /kernel/drv/lpfc.conf file to loop (topology=4).	The topology option in the /kernel/drv/lpfc.conf file is set to the proper value automatically when you configure the host. Use this option to reset the topology only if you change your configuration after you install.

Option	Description	Notes
-m <cfmode>	Specifies the cfmode. Arguments are: ◆ single_image ◆ partner ◆ dual_fabric ◆ standby ◆ mixed ◆ stand_alone	To determine the cfmode, perform the following steps: 1. Connect to the storage system. You can use the telnet command to do this. 2. Execute the fcp show cfmode command.
-p	Prints out the existing settings.	~
-r	Reverts changes to the /kernel/drv/lpfc.conf and /etc/system files made by the basic_config script. This is done by copying the backup file to the current file.	Important: The basic_config script allows you to revert only one time. To revert more than one time, you must create your own backup copy of the /kernel/drv/lpfc.conf and /etc/system file.
-s	For switched configurations, sets the topology parameter in the /kernel/drv/lpfc.conf file to switch point-to-point (topology=2).	The topology option in the /kernel/drv/lpfc.conf file is set to the proper value automatically when you configure the host. Use this option to reset the topology only if you change your configuration after you install.
-v	Displays version information for the basic_config script.	~

Example of using basic_config on systems with LPFC drivers

The following example steps you through the process of using basic_config to configure system parameters on a system that is using Emulex LPFC HBAs.

Step	Action
1	Check the current values for the parameters. Enter the following: basic_config -p Result: The software displays the current setting for NetApp required items in the /etc/system and /kernel/drv/lpfc.conf file.
2	Use the basic_config -m command line to configure /etc/system and /kernel/drv/lpfc.conf file items to the recommended setting. The sample command line shown below uses a cfmode of single image. basic_config -m single_image Result: The software displays recommended settings and prompts you for confirmation. You must enter y at the prompt that appears to confirm the changes and update the files. Example: This example uses the basic_config -m single_image command. The command displays the variables it will reset. Then it asks you to confirm this action. <pre>basic_config -m single_image Verified Emulex driver is loaded. Continuing configuration... /etc/system file will have these changes: set sd:sd_io_time=0x78 Emulex lpfc.conf will have these changes: lun-queue-depth=16; tgt-queue-depth=256; scan-down=0; num-iocbs=1024; num-bufs=1024; dqfull-throttle-up-time=4; dqfull-throttle-up-inc=2; nodev-tmo=120; linkdown-tmo=30; Please confirm the changes? (y/yes/n/no) y Please reboot your host for the changes to take effect.</pre>

Configuring parameters on systems using Sun native drivers

Tasks involved in configuring systems using native drivers

If your system uses Sun native drivers, you must complete the configuration by setting the `/kernel/drv/ssd.conf` parameters to required values. There are two ways to set this values:

- ◆ Manually enter the values by editing the file.
- ◆ Automatically set the values by running the `basic_config` command.

The sections that follow describe these tasks.

Setting `ssd.conf` variables manually on systems using native drivers

If your system uses Sun native drivers, you need to modify the `/kernel/drv/ssd.conf` file. The required values are:

Note

Versions of the Host Utilities using native drivers always use single image `cfmode`. If you are using native drivers, verify that you are also using single image `cfmode`.

- ◆ `throttle_max=8`
- ◆ `not_ready_retries=300`
- ◆ `busy_retries=30`
- ◆ `reset_retries=30`
- ◆ `throttle_min=2`

The Solaris Host Utilities provides a best-fit setting for target and/or LUN queue depths. NetApp provides documentation for determining host queue depths and sizing. For information on NetApp FC and iSCSI storage systems, see:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

For information on configuration limitations associated with NetApp storage controllers, see the *iSCSI and Fibre Channel Configuration Guide*:

http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/config_guide/setting_up/concept/c_oc_set_config-limits-overview.html

Setting VIP/PID information

You must also set the VIP/PID information to “NETAPP LUN”.

You can use the `basic_config` command to configure these parameters. It uses the following bit mask settings.

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";  
netapp-ssd-config=1,0x9007,8,300,30,0,0,0,0,0,0,0,0,0,30,0,0,2,0,0;
```

Using `basic_config` to set the options on systems using native drivers and Veritas 5.0

On systems using Sun native drivers and Veritas Storage Foundation 5.0 DMP, the `basic_config` command has the following format:

```
basic_config [-vxdmp_set |  
-vxdmp_unset | -vxdmp_set_force | -vxdmp_unset_force]
```

The following table explains these options. For an example of using the `basic_config` command, see [“Example of using `basic\_config` with native drivers”](#) on page 64.

Option	Description	Notes
-h, -help	Displays the usage page.	~
-p	Prints out the existing settings.	~
-v	Displays version information for the <code>basic_config</code> script.	~
-vxdmp_set	Adds the required settings to <code>ssd.conf</code> for VxDMP with native drivers.	For systems using Veritas DMP.
-vxdmp_set_force	Performs the same tasks as the <code>-vxdmp_set</code> option, but suppresses the confirmation prompt.	For systems using Veritas DMP.
-vxdmp_unset	Removes the entries added by the <code>-vxdmp_set</code> option.	For systems using Veritas DMP. Important: The <code>basic_config</code> script allows you to revert only one time. To revert more than one time, you must create your own backup copy of the <code>/kernel/drv/ssd.conf</code> file.

Option	Description	Notes
-vxdmp_unset_force	Performs the same tasks as the -vxdmp_unset option, but suppresses the confirmation prompt.	For systems using DMP.

Example of using basic_config with native drivers

The following example illustrates how to using basic_config to configure system parameters on a system that is using Sun native drivers.

Step	Action
1	Use the basic_config -vxdmp_set command to configure the appropriate parameters in the /kernel/drv/ssd.conf file to the recommended values. Example: The following example shows the type of output produced by the basic_config -vxdmp_set command when you execute it on a SPARC system. <pre># /opt/NTAP/SANToolkit/bin/basic_config -vxdmp_set W A R N I N G This script will modify /kernel/drv/ssd.conf to add settings required for your storage system. Do you wish to continue (y/n)?-->y The original version of the /kernel/drv/ssd.conf file has been saved to /kernel/drv/ssd.conf.1228864383 Please reboot now</pre>

Chapter overview

This chapter explains how to configure the Solaris host for the MPxIO stack. It provides information on the system parameters that need to be changed and methods for changing them.

Topics in this chapter

This chapter includes the following topic:

- ◆ [“Completing the installation and configuring system parameters”](#) on page 66

Completing the installation and configuring system parameters

Completing the Host Utilities setup on a MPxIO stack

To complete the configuration when you're using a MPxIO stack, you must modify the parameters in either `/kernel/drv/ssd.conf` or `/kernel/drv/sd.conf` and set them to the recommended values. To set the recommended values, you can either

- ◆ Manually edit the file for your system.
- ◆ Use the `basic_config` command to automatically make the changes.

Note

For a complete list of the host parameters that NetApp recommends you change and an explanation of why that change is recommended, see the *Host Settings Affected by the Host Utilities* document, which is on the NOW site on the SAN/iSAN Information Library page (<http://now.netapp.com/NOW/knowledge/docs/san/>).

Before you configure the system parameters

Complete the following tasks before you configure system parameters.

- ◆ Create your own backup of the file you are modifying:
 - ❖ `/kernel/drv/ssd.conf` for systems using SPARC processors
 - ❖ `/kernel/drv/sd.conf` for systems using x86/x64 (AMD) processors
- The `basic_config` command automatically creates backups for you, but you can only revert to those backups once. By manually creating the backups, you can revert to them as needed.
- ◆ If MPxIO was previously installed using Host Utilities or Host Attach Kit prior to 3.0.1 and ALUA was not enabled, you must remove it.

Enable ALUA for MPxIO

If you had installed a version of the Host Utilities prior to 3.0.1 and were not using ALUA, you must enable ALUA. To do this, execute the command:

```
# mpxio_set -d
```

This script removes the symmetric options from the `scsi_vhci.conf` file.

You must reboot your system for these changes to take effect.

Manually changing parameters for systems using MPxIO

You can manually modify the `/kernel/drv/ssd.conf` file (SPARC processor) or the `/kernel/drv/sd.conf` file (x86/64 (AMD) processor). The required values are the same for both SPARC processor systems and x86/x64 (AMD) processor systems using MPxIO:

- ◆ `throttle_max=64`
- ◆ `not_ready_retries=300`
- ◆ `busy_retries=30`
- ◆ `reset_retries=30`
- ◆ `throttle_min=8`

You must also set the VIP/PID information to “NETAPP LUN”. You can use the `basic_config` command to configure this information.

Using `basic_config` to configure system parameters for MPxIO stacks

The `basic_config` command is a configuration script included in the Host Utilities SAN Toolkit that you can use to automatically set system values to the ones recommended for use with the Host Utilities, including setting the VIP/PID. This command updates the items in `/kernel/drv/ssd.conf` or `/kernel/drv/sd.conf` to the values recommended for NetApp storage systems using Sun native drivers.

The arguments you use with the `basic_config` command vary depending on whether you are using a SPARC processor system or an x86/x64 (AMD) processor system.

For SPARC processor systems: Run the `basic_config -ssd_set` command. This command uses a bit mask to modify the `/kernel/drv/ssd.conf` to add the settings required for your storage system.

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";  
netapp-ssd-config=1,0x9007,8,300,30,0,0,0,0,0,0,0,0,0,30,0,0,2,0,0;
```

```
ssd-config-list="NETAPP LUN", "netapp-ssd-config";  
netapp-ssd-config=1,0x9007,64,300,30,0,0,0,0,0,0,0,0,0,30,0,0,8,0,0;
```

For x86/x64 (AMD) processor systems: Run the `basic_config -sd_set` command. This command uses a bit mask to modify the `/kernel/drv/sd.conf` to add the settings required for your storage system.

```
sd-config-list="NETAPP LUN", "netapp-sd-config";  
netapp-sd-config=1,0x9c01,64,0,0,0,0,0,0,0,0,0,0,300,30,30,0,0,8,0,0;
```

Options available with `basic_config`

When using Sun native drivers with MPxIO, the `basic_config` command accepts the following options.

Option	Description	Notes
-h, -help	Displays the usage page.	~
-p	Prints out the existing settings.	~
-ssd_set	Adds the required settings to <code>sd.conf</code> or <code>ssd.conf</code> for Sun native drivers.	For SPARC systems using MPxIO.
-ssd_set_force	Performs the same tasks as the <code>-ssd_set</code> option, but suppresses the confirmation prompt.	For SPARC systems using MPxIO.
-ssd_unset	Removes the entries added by the <code>-ssd_set</code> option.	For SPARC systems using MPxIO.
-ssd_unset_force	Performs the same tasks as the <code>-ssd_unset</code> option, but suppresses the confirmation prompt.	For SPARC systems using MPxIO.
-sd_set	Adds the required settings to <code>sd.conf</code> or <code>sd.conf</code> for Sun native drivers.	For x86/x64 (AMD) systems using MPxIO.
-sd_set_force	Performs the same tasks as the <code>-sd_set</code> option, but suppresses the confirmation prompt.	For x86/x64 (AMD) systems using MPxIO.
-sd_unset	Removes the entries added by the <code>-sd_set</code> option.	For x86/x64 (AMD) systems using MPxIO.
-sd_unset_force	Performs the same tasks as the <code>-sd_unset</code> option, but suppresses the confirmation prompt.	For x86/x64 (AMD) systems using MPxIO.
-sd_help	Displays the <code>sd_config</code> help page.	~
-v	Displays version information for the <code>basic_config</code> script.	~

**Example of
changing the
parameters using
basic_config**

To change the parameters using the basic_config command system parameters, complete the steps below.

Note_____

If you need to remove these changes, run the basic_config -ssd_unset command on SPARC systems and the basic_config -sd_unset command on x86/x64 (AMD) systems.

Step	Action
1	Execute the appropriate basic_config command for your system. The basic_config command has the following format: ◆ For SPARC systems - basic_config [-ssd_set -ssd_unset -ssd_set_force -ssd_unset_force] ◆ For x86/x64 (AMD) systems, execute - basic_config [-sd_set -sd_unset -sd_set_force -sd_unset_force] Example: The following example shows the type of output produced by the basic_config -ssd_set command when you execute it on a SPARC system. <pre># /opt/NTAP/SANToolkit/bin/basic_config -ssd_set W A R N I N G This script will modify /kernel/drv/ssd.conf to add settings required for your storage system. Do you wish to continue (y/n)?--->y The original version of /kernel/drv/ssd.conf has been saved to /kernel/drv/ssd.conf.1158267336 Please reboot now</pre>

Step	Action
2	To enable MPxIO on FC systems, do the following: ◆ For SPARC systems running Solaris 10 update 5 and later, enter: <pre># stmsboot -D fp -e</pre> ◆ For SPARC systems running Solaris 10 releases prior to Update 5 enter: <pre># stmsboot -e</pre> Note _____ This command is also available for x86/x64 systems running Solaris 10 update 4. ◆ For x86/x64 (AMD) systems, MPxIO is on by default. You need to manually reboot the host to enable the FC parameters.
3	You must reboot your system for these changes to take effect. Example: You can use the following command to reboot your system. <pre># touch /reconfigure # init 6</pre>

Chapter overview

The Veritas DMP stack requires that you install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM) for NetApp storage systems. This chapter provides information on installing and configuring the ASL and APM and how these components affect multipathing and failover.

Topics in this chapter

Topics in this appendix include the following:

- ◆ [“Overview of the Array Support Library and Array Policy Module”](#) on page 72
- ◆ [“Installing or Upgrading the ASL and the APM”](#) on page 74
- ◆ [“Example of uninstalling and installing the ASL and APM”](#) on page 79
- ◆ [“Obtaining multipathing information”](#) on page 83
- ◆ [“Optimizing the Veritas DMP Restore Daemon”](#) on page 88
- ◆ [“Understanding ASL error message types”](#) on page 90

Overview of the Array Support Library and Array Policy Module

The Array Support Library and the Array Policy Module

Symantec provides both the Array Support Library (ASL) and the Array Policy Module (APM). You download them from the Symantec Web site.

To determine which version of the ASL and APM you need for this version of the Host Utilities, check the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>. This information is updated frequently.

Note

Because these are Symantec products, Symantec provides customer support if you encounter a problem using them.

What the ASL is

The ASL is a NetApp-qualified library provides information about storage array attributes configurations to the Device Discovery Layer (DDL) of VxVM.

The DDL is a component of VxVM that discovers available enclosure information for disks and disk arrays that are connected to a host system. The DDL calls ASL functions during the storage discovery process on the host. The ASL in turn “claims” a device based on vendor and product identifiers. The claim associates the storage array model and product identifiers with the device.

You can add or remove ASLs from a running VxVM system. You do not need to reboot the host.

Note

In a Veritas Storage Foundation RAC cluster, you must stop clustering on a node before you remove the ASL.

What the APM is

The APM is a kernel module that defines I/O error handling, failover path selection, and other failover behavior for a specific array. The NetApp APM is customized to optimize I/O error handling and failover path selection for the NetApp environment.

What the ASL provides

The ASL provides enclosure-based naming information about SAN-attached storage systems. With enclosure-based naming, the name of the Veritas disk contains the model name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to VxVM about SAN-attached storage systems, instead of referring to them as JBOD devices or raw devices.

By using the ASL, you obtain the following information about the LUNs:

- ◆ Enclosure name—VxVM's enclosure-based naming feature creates disk names based on the name of its enclosure, or disk array, and not a raw device name.
- ◆ Multipathing policy—Whether the NetApp storage is accessed as an Active/Active (A/A-NETAPP) disk array or an Active/Passive Concurrent (A/P-C-NETAPP) disk array. The ASL also provides information about primary and secondary paths to NetApp storage.

Using storage with the ASL

You cannot use storage systems simultaneously as JBODs and vendor arrays. If you install the ASL, storage systems cannot be configured in VxVM as JBODs. They are reported as storage arrays unless you explicitly exclude them by using the `vxddladm exclude array` command.

Installing or Upgrading the ASL and the APM

Upgrading to a new ASL and APM when running Storage Foundation 5.0

If you are using DMP with Veritas Storage Foundation 5.0 or higher, you need to install the ASL and the APM. If you have these installed and need to upgrade, you must first remove the older versions. You can use the `pkgrm` command to uninstall the ASL and APM. For information on removing them, see [“Uninstalling the ASL and the APM”](#) on page 79.

ASL and APM installation overview

Installing the ASL and the APM involves the following tasks:

- ◆ Verifying that your configuration meets system requirements. See the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do> for current information about the system requirements.
- ◆ Determining the ASL version if it is currently installed.
- ◆ Obtaining the ASL and the APM.
- ◆ Installing a new version of the ASL and APM or upgrading an existing version of the ASL and APM.

Determining the ASL version

If you currently have the ASL installed, use the Veritas `vxddladm listversion` command to determine its version.

Example: The following is output from the `vxddladm listversion` command.

```
# vxddladm listversion
```

LIB_NAME	ASL_VERSION	Min. VXVM version
libvxCLARiiON.so	vm-5.0-rev-1	5.0
libvxcscovrts.so	vm-5.0-rev-1	5.0
libvxemc.so	vm-5.0-rev-2	5.0
libvxengenio.so	vm-5.0-rev-1	5.0
libvxhds9980.so	vm-5.0-rev-1	5.0
libvxhdsalua.so	vm-5.0-rev-1	5.0
libvxhdsusp.so	vm-5.0-rev-2	5.0
libvxhpalua.so	vm-5.0-rev-1	5.0
libvxibmds4k.so	vm-5.0-rev-1	5.0
libvxibmds6k.so	vm-5.0-rev-1	5.0
libvxibmds8k.so	vm-5.0-rev-1	5.0
libvxsenas.so	vm-5.0-rev-1	5.0

libvxshark.so	vm-5.0-rev-1	5.0
libvxsunse3k.so	vm-5.0-rev-1	5.0
libvxsunset4.so	vm-5.0-rev-1	5.0
libvxvpath.so	vm-5.0-rev-1	5.0
libvxp1281024.so	vm-5.0-rev-1	5.0
libvxp12k.so	vm-5.0-rev-2	5.0
libvxibmsvc.so	vm-5.0-rev-1	5.0
libvxnetapp.so	vm-5.0-rev-0	5.0

Obtaining the ASL and APM software

The ASL and APM are not distributed with the Host Utilities. You obtain the ASL and APM from the Symantec web site. For Veritas Storage Foundation 5.0, the Symantec TechNote download file contains the software packages for both the ASL and the APM. You must extract these software packages and then install each one separately as described in the TechNote.

For information on the link to the ASL/APM TechNote on the Symantec Web site, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Installing the ASL and APM software you downloaded

To install a fresh version of the ASL and APM that you downloaded from Symantec, complete the following steps.

Note

Before you begin, make sure you obtain the ASL/APM TechNote, which you can review at the Symantec Web site. The TechNote contains Symantec's instructions for installing the NetApp ASL and APM.

Step	Action
1	Log in to the VxVM system as the root user.
2	If you already have NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering: vxddladm rmjbod vid=NETAPP
3	Check the NetApp Interoperability Matrix at http://now.netapp.com/matrix/mtx/login.do to determine the currently supported version of the ASL and APM. Follow the link in the matrix to the ASL/APM TechNote on the Symantec Web site.

Step	Action
4	Install the ASL and APM according to the installation instructions provided by the ASL/APM TechNote on the Symantec Web site. Note _____ You should have your LUNs set up before you install the ASL and APM. For information on setting up LUNs, see “Basic Tasks for Working with LUNs” on page 91.
5	If your host is connected to NetApp storage, you can verify the installation by locating the NetApp Enclosure Type in the output of the <code>vxddmpadm listenclosure all</code> command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. vxddmpadm listenclosure all Example: The <code>vxddmpadm listenclosure all</code> command shows the Enclosure Types as FAS3020. To make this example easier to read, the line for the NetApp storage is shown in bold <pre># vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE ===== FAS30200 FAS3020 1071852 CONNECTED A/P-C-NETAPP Disk Disk DISKS CONNECTED Disk</pre>

Step	Action
6	If your host is not connected to storage, use the following command: <pre>vxddladm listsupport all</pre> Example: The following is a sample of the output you see when you enter the vxddladm listsupport all command. To make this example easier to read, the line for the NetApp storage is shown in bold. <pre># vxddladm listsupport all LIBNAME VID ===== libvxCLARiiON.so DGC libvxcscovrts.so CSCOVRTS libvxemc.so EMC libvxengenio.so SUN libvxhds9980.so HITACHI libvxhdsalua.so HITACHI libvxhdsusp.so HITACHI libvxhpalua.so HP, COMPAQ libvxibmds4k.so IBM libvxibmds6k.so IBM libvxibmds8k.so IBM libvxsenas.so SENA libvxshark.so IBM libvxsunse3k.so SUN libvxsunset4.so SUN libvxvpath.so IBM libvxxp1281024.so HP libvxxp12k.so HP libvxibmsvc.so IBM libvxnetapp.so NETAPP</pre>

Step	Action																																																																																																				
7	Enter the <code>vxddmpadm listapm all</code> command to verify that the APM is installed. Example: This command produces the following type of output. To make this example easier to read, the line for the NetApp storage is shown in bold <table><tr><th>Filename</th><th>APM Name</th><th>APM Version</th><th>Array Types</th><th>State</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>dmpaa</td><td>dmpaa</td><td>1</td><td>A/A</td><td>Active</td></tr><tr><td>dmpaaa</td><td>dmpaaa</td><td>1</td><td>A/A-A</td><td>Not-Active</td></tr><tr><td>dmpsvc</td><td>dmpsvc</td><td>1</td><td>A/A-IBMSVC</td><td>Not-Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P-C</td><td>Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-VERITAS</td><td>Not-Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-T3PLUS</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG-C</td><td>Not-Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>Disk</td><td>Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>APdisk</td><td>Active</td></tr><tr><td>dmpahsalua</td><td>dmpahsalua</td><td>1</td><td>A/A-A-HDS</td><td>Not-Active</td></tr><tr><td>dmpCLARiion</td><td>dmpCLARiion</td><td>1</td><td>CLR-A/P</td><td>Not-Active</td></tr><tr><td>dmpCLARiion</td><td>dmpCLARiion</td><td>1</td><td>CLR-A/PF</td><td>Not-Active</td></tr><tr><td>dmpahpalua</td><td>dmpahpalua</td><td>1</td><td>A/A-A-HP</td><td>Not-Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/A-NETAPP</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/P-C-NETAPP</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/P-NETAPP</td><td>Active</td></tr></table>	Filename	APM Name	APM Version	Array Types	State	=====	=====	=====	=====	=====	dmpaa	dmpaa	1	A/A	Active	dmpaaa	dmpaaa	1	A/A-A	Not-Active	dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active	dmpap	dmpap	1	A/P	Active	dmpap	dmpap	1	A/P-C	Active	dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active	dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active	dmpapg	dmpapg	1	A/PG	Not-Active	dmpapg	dmpapg	1	A/PG-C	Not-Active	dmpjbod	dmpjbod	1	Disk	Active	dmpjbod	dmpjbod	1	APdisk	Active	dmpahsalua	dmpahsalua	1	A/A-A-HDS	Not-Active	dmpCLARiion	dmpCLARiion	1	CLR-A/P	Not-Active	dmpCLARiion	dmpCLARiion	1	CLR-A/PF	Not-Active	dmpahpalua	dmpahpalua	1	A/A-A-HP	Not-Active	dmpnetapp	dmpnetapp	1	A/A-NETAPP	Active	dmpnetapp	dmpnetapp	1	A/P-C-NETAPP	Active	dmpnetapp	dmpnetapp	1	A/P-NETAPP	Active
Filename	APM Name	APM Version	Array Types	State																																																																																																	
=====	=====	=====	=====	=====																																																																																																	
dmpaa	dmpaa	1	A/A	Active																																																																																																	
dmpaaa	dmpaaa	1	A/A-A	Not-Active																																																																																																	
dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active																																																																																																	
dmpap	dmpap	1	A/P	Active																																																																																																	
dmpap	dmpap	1	A/P-C	Active																																																																																																	
dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active																																																																																																	
dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active																																																																																																	
dmpapg	dmpapg	1	A/PG	Not-Active																																																																																																	
dmpapg	dmpapg	1	A/PG-C	Not-Active																																																																																																	
dmpjbod	dmpjbod	1	Disk	Active																																																																																																	
dmpjbod	dmpjbod	1	APdisk	Active																																																																																																	
dmpahsalua	dmpahsalua	1	A/A-A-HDS	Not-Active																																																																																																	
dmpCLARiion	dmpCLARiion	1	CLR-A/P	Not-Active																																																																																																	
dmpCLARiion	dmpCLARiion	1	CLR-A/PF	Not-Active																																																																																																	
dmpahpalua	dmpahpalua	1	A/A-A-HP	Not-Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/A-NETAPP	Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/P-C-NETAPP	Active																																																																																																	
dmpnetapp	dmpnetapp	1	A/P-NETAPP	Active																																																																																																	

Next steps

After you install the ASL and APM, complete the following procedures:

- ◆ If you have Data ONTAP 7.1 or higher, it is recommended to change the `cfmode` setting of your clustered systems to single image and then reconfigure your host to discover the new paths to the disk. See [“Where to go for more information”](#) on page 15.
- ◆ On the storage system, create LUNs and map them to igroups containing the WWPNs of the host HBAs. See [“Creating and mapping igroups and LUNs”](#) on page 96.
- ◆ On the host, discover the new LUNs and configure them to be managed by VxVM. See [“Basic Tasks for Working with LUNs”](#) on page 91.

Example of uninstalling and installing the ASL and APM

Before you begin

You should not uninstall the ASL and APM when you have I/O running. Before you uninstall the ASL and APM, make sure you:

- ◆ Quiesce I/O.
- ◆ Deport the disk group.

Uninstalling the ASL and the APM

The following is a sample uninstall of the ASL and the APM. If you were actually doing this, your output would vary slightly based on your system setup.

```
# pkginfo | grep VRTSNTAP
system      VRTSNTAPapm          Veritas NetApp Array Policy Module.
system      VRTSNTAPasl          Veritas NetApp Array Support Library
# pkgrm VRTSNTAPapm
```

The following package is currently installed:

```
VRTSNTAPapm  Veritas NetApp Array Policy Module.
              (sparc) 5.0,REV=09.12.2007.16.16
```

Do you want to remove this package? [y,n,?,q] y

```
## Removing installed package instance <VRTSNTAPapm>
```

This package contains scripts which will be executed with super-user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y

```
## Verifying package <VRTSNTAPapm> dependencies in global zone
```

```
## Processing package information.
```

```
## Executing preremove script.
```

```
Check if Module is loaded
```

```
## Removing pathnames in class <none>
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm/sparcv9 <shared pathname not removed>
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
```

```
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
```

```
/kernel/drv/vxapm <shared pathname not removed>
```

```

/kernel/drv <shared pathname not removed>
/kernel <shared pathname not removed>
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/64 <shared pathname not removed>
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/32 <shared pathname not removed>
/etc/vx/apmkey.d <shared pathname not removed>
/etc/vx <shared pathname not removed>
/etc <shared pathname not removed>
## Updating system information.

Removal of <VRTSNTAPapm> was successful.
#

# pkgrm VRTSNTAPasl

The following package is currently installed:
    VRTSNTAPasl  Veritas NetApp Array Support Library
                  (sparc) 5.0,REV=11.19.2007.14.03

Do you want to remove this package? [y,n,?,q] y

## Removing installed package instance <VRTSNTAPasl>

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package <VRTSNTAPasl> dependencies in global zone
## Processing package information.
## Executing preremove script.
Unloading the library
## Removing pathnames in class <none>
/etc/vx/lib/discovery.d/libvxnetapp.so.2
/etc/vx/lib/discovery.d <shared pathname not removed>
/etc/vx/lib <shared pathname not removed>
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/aslkey.d <shared pathname not removed>
## Executing postremove script.
## Updating system information.

Removal of <VRTSNTAPasl> was successful.
#

```

Installing the ASL and the APM

The following is a sample installation of the ASL and the APM when you have Veritas Storage Foundation 5.0. If you were actually doing this, your output would vary slightly based on your system setup.

```
# pkgadd -d . VRTSNTAPas1

Processing package instance <VRTSNTAPas1> from </tmp>

Veritas NetApp Array Support Library(sparc) 5.0,REV=11.19.2007.14.03
Copyright © 1990-2006 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered trademarks of
Symantec Corporation or its affiliates in the U.S. and other countries. Other
names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer
software" and "commercial computer software documentation" as defined in FAR
Sections 12.212 and DFARS Section 227.7202.
Using </etc/vx> as the package base directory.
## Processing package information.
## Processing system information.
   3 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <VRTSNTAPas1> [y,n,?] y

Installing Veritas NetApp Array Support Library as <VRTSNTAPas1>

## Installing part 1 of 1.
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/lib/discovery.d/libvxnetapp.so.2
[ verifying class <none> ]
## Executing postinstall script.
Adding the entry in supported arrays
Loading The Library
```

```

Installation of <VRTSNTAPasl> was successful.
#

# pkgadd -d . VRTSNTAPapm

Processing package instance <VRTSNTAPapm> from </tmp>
Veritas NetApp Array Policy Module.(sparc) 5.0,REV=09.12.2007.16.16
    Copyright 1996-2005 VERITAS Software Corporation. All rights reserved.
    VERITAS, VERITAS SOFTWARE, the VERITAS logo and all other VERITAS
    product names and slogans are trademarks or registered
    trademarks of VERITAS Software Corporation in the USA and/or other
    countries. Other product names and/or slogans mentioned herein may
    be trademarks or registered trademarks of their respective companies.

Using </> as the package base directory.
## Processing package information.
## Processing system information.
    9 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of <VRTSNTAPapm> [y,n,?] y

Installing Veritas NetApp Array Policy Module. as <VRTSNTAPapm>

## Installing part 1 of 1.
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
[ verifying class <none> ]
## Executing postinstall script.

Installation of <VRTSNTAPapm> was successful.
#

```

Obtaining multipathing information

What an ASL array type is

The ASL reports information about the multipathing configuration to the DDL as one of the following disk array types:

- ◆ Active/Active NetApp (A/A-NETAPP)—All paths to storage are active and simultaneous I/O is supported on all paths. If a path fails, I/O is distributed across the remaining paths.
- ◆ Active/Passive Concurrent-NetApp (A/P-C-NETAPP) - The array supports concurrent I/O and load balancing by having multiple primary paths to LUNs. Failover to the secondary (passive) path occurs only if all the active primary paths fail.

For additional information, see the *Veritas Volume Manager Administrator's Guide* for details about system management.

How the storage system's FC failover mode corresponds to array types

In clustered NetApp storage configurations, the array type corresponds to the storage system cfmode settings. If you use the standby cfmode, the array type will be A/A-NETAPP; otherwise, it will be A/P-C-NETAPP.

Note

The ASL also supports direct attached, non-clustered configurations, including NearStore models. These configurations have no cfmode settings. ASL reports these configurations as Active/Active (A/A-NETAPP) array types.

How sanlun displays the array type

The Host Utilities SAN Toolkit includes the sanlun utility in the opt/NTAP/SANToolkit/bin directory. You use the sanlun utility to display information about paths to LUNs on the storage system.

- ◆ When the ASL is installed and the LUN is controlled by VxVM, the output of the sanlun command displays the Multipath_Policy as A/P-C or A/A.
- ◆ If the LUN is not controlled by a volume manager, then the Multipathing_Policy is none and the Multipathing-provider is none.

Using VxVM to display available paths

To obtain information about available paths to a LUN using VxVM, complete the following steps

Step	Action
1	Enter the following command to view all devices: <code>vxdisk list</code> Result: The VxVM management interface displays the vxdisk device, type, disk, group and status. It also shows which disks are managed by VxVM. Example: The following is an example of the type of output you see with the vxdisk list command. <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:none - - online invalid Disk_1 auto:none - - online invalid FAS30201_0 auto:cdsdisk c2t4d60s2 dg2 online shared FAS30201_1 auto:cdsdisk c2t4d26s2 dg1 online shared FAS30201_2 auto:cdsdisk - - online FAS30201_3 auto:cdsdisk c2t4d61s2 dg2 online shared FAS30201_4 auto:cdsdisk c2t4d84s2 dg2 online shared</pre> (The output has been truncated to make the document easier to read.)

Step	Action
2	On the host console, enter the following command to display path information for the device you want: <pre>vxddmpadm getsubpaths dmpnodename=<device></pre> where <i><device></i> is the name listed under the output of the <code>vxdisk list</code> command. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths dmpnodename=FAS30201_0 NAME STATE[A] PATH-TYPE[M] CTLR-NAME ENCLR-TYPE ENCLR-NAME ATTRS ===== = c2t4d60s2 ENABLED(A) PRIMARY c2 FAS3020 FAS30201 - c2t5d60s2 ENABLED(A) PRIMARY c2 FAS3020 FAS30201 - c2t6d60s2 ENABLED SECONDARY c2 FAS3020 FAS30201 - c2t7d60s2 ENABLED SECONDARY c2 FAS3020 FAS30201 - c3t4d60s2 ENABLED(A) PRIMARY c3 FAS3020 FAS30201 - c3t5d60s2 ENABLED(A) PRIMARY c3 FAS3020 FAS30201 - c3t6d60s2 ENABLED SECONDARY c3 FAS3020 FAS30201 - c3t7d60s2 ENABLED SECONDARY c3 FAS3020 FAS30201 -</pre>

Step	Action																																																																																											
3	Enter the following command to obtain path information for a host HBA: <pre>vxddmpadm getsubpaths ctlr=<controller_name></pre> controller_name is the controller displayed under “CTRL-NAME” in the output of the vxddmpadm getsubpaths dmpnodename command you entered in Step 2. Result: The output displays information about the paths to the filer storage (whether the path is a primary or secondary path). The output also lists the filer that the device is mapped to. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths ctlr=c2</pre> <table><tr><th>NAME</th><th>STATE[A]</th><th>PATH-TYPE[M]</th><th>DMPNODENAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td colspan="7">=====</td></tr><tr><td>=</td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>c2t4d0s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_81</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d1s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_80</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d2s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_76</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d3s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_44</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d4s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_45</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d5s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_2</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d6s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_82</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d7s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_79</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d8s2</td><td>ENABLED(A)</td><td>PRIMARY</td><td>FAS30201_43</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr><tr><td>c2t4d9s2</td><td>ENABLED</td><td>SECONDARY</td><td>FAS30201_41</td><td>FAS3020</td><td>FAS30201</td><td>-</td></tr></table> (The output has been truncated to make the document easier to read.)	NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====							=							c2t4d0s2	ENABLED(A)	PRIMARY	FAS30201_81	FAS3020	FAS30201	-	c2t4d1s2	ENABLED	SECONDARY	FAS30201_80	FAS3020	FAS30201	-	c2t4d2s2	ENABLED(A)	PRIMARY	FAS30201_76	FAS3020	FAS30201	-	c2t4d3s2	ENABLED	SECONDARY	FAS30201_44	FAS3020	FAS30201	-	c2t4d4s2	ENABLED(A)	PRIMARY	FAS30201_45	FAS3020	FAS30201	-	c2t4d5s2	ENABLED	SECONDARY	FAS30201_2	FAS3020	FAS30201	-	c2t4d6s2	ENABLED(A)	PRIMARY	FAS30201_82	FAS3020	FAS30201	-	c2t4d7s2	ENABLED	SECONDARY	FAS30201_79	FAS3020	FAS30201	-	c2t4d8s2	ENABLED(A)	PRIMARY	FAS30201_43	FAS3020	FAS30201	-	c2t4d9s2	ENABLED	SECONDARY	FAS30201_41	FAS3020	FAS30201	-
NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																																						
=====																																																																																												
=																																																																																												
c2t4d0s2	ENABLED(A)	PRIMARY	FAS30201_81	FAS3020	FAS30201	-																																																																																						
c2t4d1s2	ENABLED	SECONDARY	FAS30201_80	FAS3020	FAS30201	-																																																																																						
c2t4d2s2	ENABLED(A)	PRIMARY	FAS30201_76	FAS3020	FAS30201	-																																																																																						
c2t4d3s2	ENABLED	SECONDARY	FAS30201_44	FAS3020	FAS30201	-																																																																																						
c2t4d4s2	ENABLED(A)	PRIMARY	FAS30201_45	FAS3020	FAS30201	-																																																																																						
c2t4d5s2	ENABLED	SECONDARY	FAS30201_2	FAS3020	FAS30201	-																																																																																						
c2t4d6s2	ENABLED(A)	PRIMARY	FAS30201_82	FAS3020	FAS30201	-																																																																																						
c2t4d7s2	ENABLED	SECONDARY	FAS30201_79	FAS3020	FAS30201	-																																																																																						
c2t4d8s2	ENABLED(A)	PRIMARY	FAS30201_43	FAS3020	FAS30201	-																																																																																						
c2t4d9s2	ENABLED	SECONDARY	FAS30201_41	FAS3020	FAS30201	-																																																																																						

Using sanlun to obtain multipathing information

To obtain multipathing information by using the sanlun command, complete the following step.

Step	Action																																																															
1	On the host command line, enter the following command: <pre>sanlun lun show -p all</pre> Result: sanlun displays path information for each LUN. For LUNs managed by VxVM, the output displays the Veritas multipathing policy (A/A or A/PG-C array types) and the path I/O policy. Example: The following partial output shows eight paths to one LUN managed by VxVM. <pre># sanlun lun show -p all filerX:/vol/vol1/data_lun60 (LUN 60) /dev/rdsk/c2t4d60s2 (FAS30201_0) 9g (9663676416) lun state: GOOD</pre> <table><tr><th>path state</th><th>path type</th><th>device filename</th><th>host HBA</th><th>target port</th><th>partner port</th><th>Multipath policy</th></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c2t4d60s2</td><td>lpfc0</td><td>1a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c2t5d60s2</td><td>lpfc0</td><td>1b</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t5d60s2</td><td>lpfc1</td><td>4a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t4d60s2</td><td>lpfc1</td><td>4b</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c2t7d60s2</td><td>lpfc0</td><td></td><td>1b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c2t6d60s2</td><td>lpfc0</td><td></td><td>1a</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t6d60s2</td><td>lpfc1</td><td></td><td>4b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t7d60s2</td><td>lpfc1</td><td></td><td>4a</td><td>A/PG-C</td></tr></table>	path state	path type	device filename	host HBA	target port	partner port	Multipath policy	enabled	primary	/dev/rdsk/c2t4d60s2	lpfc0	1a		A/PG-C	enabled	primary	/dev/rdsk/c2t5d60s2	lpfc0	1b		A/PG-C	enabled	primary	/dev/rdsk/c3t5d60s2	lpfc1	4a		A/PG-C	enabled	primary	/dev/rdsk/c3t4d60s2	lpfc1	4b		A/PG-C	enabled	secondary	/dev/rdsk/c2t7d60s2	lpfc0		1b	A/PG-C	enabled	secondary	/dev/rdsk/c2t6d60s2	lpfc0		1a	A/PG-C	enabled	secondary	/dev/rdsk/c3t6d60s2	lpfc1		4b	A/PG-C	enabled	secondary	/dev/rdsk/c3t7d60s2	lpfc1		4a	A/PG-C
path state	path type	device filename	host HBA	target port	partner port	Multipath policy																																																										
enabled	primary	/dev/rdsk/c2t4d60s2	lpfc0	1a		A/PG-C																																																										
enabled	primary	/dev/rdsk/c2t5d60s2	lpfc0	1b		A/PG-C																																																										
enabled	primary	/dev/rdsk/c3t5d60s2	lpfc1	4a		A/PG-C																																																										
enabled	primary	/dev/rdsk/c3t4d60s2	lpfc1	4b		A/PG-C																																																										
enabled	secondary	/dev/rdsk/c2t7d60s2	lpfc0		1b	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c2t6d60s2	lpfc0		1a	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c3t6d60s2	lpfc1		4b	A/PG-C																																																										
enabled	secondary	/dev/rdsk/c3t7d60s2	lpfc1		4a	A/PG-C																																																										

Optimizing the Veritas DMP Restore Daemon

The Veritas DMP restore daemon requirements

It is recommended that you set recommended values for the restore policy of the Veritas DMP restore daemon and also specify a polling interval. These settings determine how frequently the Veritas daemon checks paths between the host and the storage system.

For the most up-to-date recommendations on the values, see the General Operational Considerations page of the *FC Solaris Host Utilities Release Notes*. The Release Notes are online in the SAN/iSAN Information Library at <http://now.netapp.com/NOW/knowledge/docs/san/>.

Setting the restore daemon interval

At the time this document was prepared, NetApp recommended you set the restore daemon interval value to 60 seconds to improve the recovery of previously failed paths.

Note_____

To see if there are new recommendations, check the Release Notes.

To set the vxrestore daemon, complete the following steps.

Step	Action
1	Stop the restore daemon. <code>/usr/sbin/vxdmpadm stop restore</code>
2	Change the restore daemon setting to 60. <code>/usr/sbin/vxdmpadm start restore interval=60 policy=check_disabled</code> Note _____

This step reconfigures and restarts the restore daemon without the need for an immediate reboot.

Step	Action
3	Edit the <code>/lib/svc/method/vxvm-sysboot</code> file in Solaris 10 or the <code>/etc/init.d/vxvm-sysboot</code> file in Solaris 9 to make the new restore daemon interval persistent across reboots. By default, the restore daemon options are: <pre>restore_daemon_opts="interval=300 policy=check_disabled"</pre> Edit the restore daemon options so that the interval is 60 seconds: <pre>restore_daemon_opts="interval=60 policy=check_disabled"</pre>
4	Save and exit the <code>/etc/init.d/vxvm-sysboot</code> file.
5	Verify the changes. <pre>/usr/sbin/vxdmpadm stat restored</pre>

Understanding ASL error message types

**Understanding
error message
severity**

The ASL works silently and seamlessly with the VxVM DDL. If an error, misconfiguration, or malfunction occurs, messages from the library are logged to the console using the host’s logging facility. The following table lists the importance and severity of these messages.

Message severity	Definition	Action required
Error	Indicates that an ERROR status is being returned from the ASL to the VxVM DDL that prevents the device (LUN) from being used. The device might still appear in vxdisk list, but it is not usable.	Call Symantec Technical Support for help.
Warning	Indicates that an UNCLAIMED status is being returned. Unless claimed by a subsequent ASL, dynamic multipathing is disabled. No error is being returned but the device (LUN) might not function as expected.	
Info	Indicates that a CLAIMED status is being returned. The device functions fully, with Veritas DMP enabled, but the results seen by the user might be other than what is expected. For example, the enclosure name might change.	

About this chapter

This chapter provides a high-level look at the tasks involved in creating and working with LUNs. It provides pointers to chapters containing instructions on configuring and managing LUNs that are specific to one Solaris environment (Veritas, or MPxIO).

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Overview of LUN configuration and management”](#) on page 92
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 96

Overview of LUN configuration and management

Task overview

LUN configuration and management involves a number of tasks. The following table summarizes the tasks for all the supported Solaris environments. If a task does not apply to all environments, then it specifies which environments it does apply to. You only need to perform the tasks that apply to your environment.

Task	Discussion	Additional information
1. Create and map igroups and LUNs	An igroup is a collection of WWPNs on the storage system that map to one or more host HBAs. After you create the igroup, you must create LUNs on the storage system, and map the LUNs to the igroup.	For complete information, refer to your version of the <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
2. (MPxIO) Enable ALUA	You must have ALUA set up to work with igroups. To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.	Refer to “(MPxIO) Check the ALUA settings on the storage system” on page 138.
3. (Veritas DMP with LPFC drivers only) Create persistent bindings	Persistent bindings permanently bind a particular target ID on the host to the storage system WWPN. You must create persistent binding between the storage system (target) and the Solaris host (initiator) to guarantee that the storage system is always available at the correct SCSI target ID on the host. You create the binding with either the <code>hbacmd</code> command or the <code>create_binding.pl</code> script.	Refer to “(LPFC drivers) Creating persistent bindings” on page 100 for information about creating bindings with <code>hbacmd</code>. Refer to “Veritas DMP: Using the <code>create_binding.pl</code> Script” on page 149 for information about creating bindings with the <code>create_binding.pl</code> script.

Task	Discussion	Additional information
4. (Veritas DMP with LPFC drivers prior to 6.21g) Modify the sd.conf file	The sd.conf file is a SCSI driver configuration file that tells the host system which disks to probe when the system is booted. You must create an entry for every NetApp LUN in the sd.conf file. You must also add entries to this file if you add additional LUNs.	Refer to “(Drivers prior to 6.21g) Modifying the sd.conf file” on page 102 for additional information.
5. (MPxIO, Sun native drivers with Veritas DMP) Display a list of controllers	If you are using an MPxIO stack or Sun native drivers with Veritas DMP, you need to get information about the controller before you can discover the LUNs. Use the <code>cfgadm -al</code> command to display a list of controllers.	For additional information, refer to the section that applies to the stack for your system uses ◆ (Sun native drivers with Veritas) “Getting the controller number” on page 105. ◆ (MPxIO) “Getting the controller number” on page 110 for additional information.

Task	Discussion	Additional information
6. Discover LUNs	(Veritas LPFC drivers) The host discovers the LUNs when you perform a reconfiguration reboot (touch /reconfigure; init 6) on the host. You can also discover LUNs using the /usr/sbin/devfsadm command. If you are using drivers 6.21g or later, you can also use the hbacmd command to discover LUNs. (MPxIO, Sun native drivers with Veritas) To discover the LUNs, use the command: <pre>/usr/sbin/cfgadm -c configure cx</pre> x is the controller number of the HBA where the LUN is expected to be visible	For additional information, refer to the section that applies to the HBAs or stack your system uses: ◆ (Veritas DMP) “(LPFC drivers) Configuring the host to discover LUNs” on page 102 ◆ (Veritas DMP) “(Sun native drivers) Configuring the host to discover LUNs” on page 105 information. ◆ (MPxIO) “Discovering LUNs” on page 110
7. Label LUNs, if appropriate for your system	Use the Solaris format utility to label the LUNs. Note— For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume. See “Specifying the LUN type for optimal performance” on page 96.	Refer to Solaris documentation for complete information about the format utility. For a quick summary, see ◆ (Veritas) “Labeling the new LUN” on page 106 ◆ (MPxIO) “Formatting or labeling the new LUN” on page 111

Task	Discussion	Additional information
8. Configure volume management software	If your configuration uses either SVM, ZFS, or VxVM, you must configure the LUNs so they are under the volume manager (SVM, VxVM) or file system (ZFS) control.	Refer to documentation for your volume manager tool for additional information. For a high-level summary, see ◆ (Veritas) “Configuring volume management” on page 108 ◆ (MPxIO) “Configuring volume management” on page 114

Creating and mapping igroups and LUNs

Tasks needed to create and map LUNs

Before you can complete the configuration, you must perform the following tasks:

- ◆ Create an igroup.

Note

If you have an active/active configuration, you must create a separate igroup on each system in the configuration.

- ◆ Create one or more LUNs and map the LUNs to an igroup.

Methods for creating igroups and LUNs

You can create igroups and LUNs on the storage system by:

- ◆ Entering the `lun setup` command.
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
- ◆ Entering a series of individual commands (such as `lun create`, `igroup create` and `lun map`).
Use this method to create one or more LUNs and igroups in any order.
- ◆ Using FilerView.
This method provides a LUN Wizard that steps you through the process of creating and mapping new LUNs.

For detailed information about creating and managing LUNs, see the *Data ONTAP® Block Access Management Guide for iSCSI and FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Specifying the LUN type for optimal performance

For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume. When you create a LUN and specify `solaris` as the value for the LUN creation `ostype` parameter, this alignment is automatic. Use the `solaris` `ostype` with UFS and VxFS file systems.

Solaris uses a newer labeling scheme, known as EFI, for LUNs that are 2TB or larger, ZFS volumes, and SVM volumes with disk sets. For these situations, you should specify `solaris_efi` as the value for the `ostype` parameter when you

create the LUN. If the `solaris_efi` ostype is not available, you must perform special steps to align the partitions to the WAFL volume. See the Solaris Host Utilities Release Notes for details.

Best practices for creating igroups and LUNs

As you create the igroups and LUNs, keep the following best practices in mind:

- ◆ Disable scheduled auto snapshots.
- ◆ Each igroup should map to an application and should include all the initiators used by the application to access its data. (Multiple applications may be using the same initiators.)
- ◆ Do not put LUNs in the storage system's root volume. The default root volume is `/vol/vol0`.

About this chapter

This chapter describes how to enable a host running Veritas Storage Foundation to discover LUNs on the storage system. It also describes how to manage LUNs on hosts utilizing Veritas and the Veritas Volume Manager (VxVM).

In some cases, the methods you use depend on whether you are using Veritas with LPFC drivers or Sun native drivers.

Note

For general information on working with LUNs, see “[Basic Tasks for Working with LUNs](#)” on page 91.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[\(LPFC drivers\) Creating persistent bindings](#)” on page 100
- ◆ “[\(LPFC drivers\) Configuring the host to discover LUNs](#)” on page 102
- ◆ “[\(Sun native drivers\) Configuring the host to discover LUNs](#)” on page 105
- ◆ “[Labeling the LUN and configuring Volume Management software](#)” on page 106

(LPFC drivers) Creating persistent bindings

What are persistent bindings?

Persistent bindings map the WWPN of each active port on the storage system to a target ID on the Solaris host initiator (HBA). Creating persistent bindings guarantees that the storage system is always available at the correct SCSI target ID on the host. Without persistent bindings, the SCSI target ID is not guaranteed.

The FC Solaris Host Utilities supports persistent bindings between a WWPN on the storage system and a target ID on the host.

Methods for creating persistent bindings

You can use any of the following tools and methods to create persistent bindings:

- ◆ **hbacmd** — This is an installation and configuration binary that loads automatically when you install the HBAnyware package. Refer to “[Creating WWPN bindings with hbacmd](#)” on page 101 for information about creating persistent bindings using hbacmd.

Note

The LPFC driver utility hbacmd is located in two different places based on your driver version. Because of this, this document does not use the full path in the examples. Please use the appropriate path based on the driver version you are installing.

- ◆ **Emulex HBAnyware** — This is the HBA management GUI provided by Emulex and included with the driver software. To configure persistent binding with HBAnyware, refer to the documentation that came with your Emulex software.
- ◆ **The create_binding.pl script** — This is a NetApp provided script that loads when you install the FC Solaris Host Utilities software.
 - ❖ To install the FC Solaris Host Utilities software, refer to “[Installing the Host Utilities software package](#)” on page 42
 - ❖ To configure persistent bindings using the create_binding.pl script, refer to “[Creating WWPN bindings with create_binding.pl](#)” on page 154.
- ◆ **Manually editing the kernel/drv/lpfc.conf file.**

You should create a backup of the /kernel/drv/lpfc.conf file before you create persistent bindings.

Creating WWPN bindings with hbacmd

Follow the steps below to create persistent bindings using hbacmd. Make a note of the output; you will use that information when you configure the host to discover LUNs.

Caution

Previous versions of the FC Solaris Host Utilities may set the automap parameter in the /kernel/drv/lpfc.conf file to 0 (automap=0). You must reset this parameter to 1 (automap=1) and reboot the host before you use hbacmd to create persistent bindings.

Step	Action
1	On the Solaris host, run hbacmd to get a list of initiators. Enter: <pre>hbacmd list</pre> Result : The software displays the list of installed adapters and their WWPNs. Write down the WWPNs. The next steps require you to enter that information.
2	Run hbacmd to print out the list of targets currently seen by the initiator. Enter: <pre>hbacmd allnodeinfo <initiator_wwpn></pre> Result : A list of targets appears as well as the SCSI bus and target information for that target. Record this information so that you can use it in the next step.
3	Run hbacmd to persistently bind the target to the initiator. Enter: <pre>hbacmd setpersistentbinding <initiator_wwpn> P P <target_wwpn> <scsi_bus> <scsi_target> <initiator_wwpn></pre> Result: The initiator will be persistently bound to the target using the target and bus numbers provided.
4	Repeat steps 2 and 3 for each target and initiator combination that you want to persistently bind.
5	If you have not set up the sd.conf file, set it up as appropriate and then reboot your system. See “(Drivers prior to 6.21g) Modifying the sd.conf file” on page 102.
6	To save the changes you have made, perform a reconfiguration reboot of the host. Enter: <pre>touch /reconfigure/sbin/init 6</pre>

(LPFC drivers) Configuring the host to discover LUNs

(Drivers prior to 6.21g) Modifying the sd.conf file

For drivers prior to 6.21g, you must add an entry for each new target ID in the sd.conf file so that the host system knows which disks to probe when the system is rebooted.

To modify the sd.conf file, complete the following steps.

Step	Action
1	Create a backup copy of the /kernel/drv/sd.conf file.
2	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for the target ID mapped to the storage system's WWPN. The entry has the following format: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> You must create an entry for every LUN mapped to the host. Note_____ If the host is connected to multiple storage systems or clusters, you must create entries in sd.conf for each additional storage system or cluster, and for each LUN mapped to the host. _____
3	On the host console, reboot the host. Enter: <pre>touch /reconfigure/sbin/init 6</pre> Note_____ If you're performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot. _____

(Drivers prior to 6.21g) Discovering LUNs

To discover new LUNs and validate that they are visible on the host with driver versions prior to 6.21g, complete the following step.

Step	Action
1	To discover new LUNs, select one of the methods below: ◆ Reboot the host with the reconfigure option: touch /reconfigure/sbin/init 6 ◆ Enter /usr/sbin/update_drv -f sd /usr/sbin/devfsadm Note _____ Occasionally the /usr/sbin/devfsadm command does not find LUNs. If this occurs, reboot the host with the reconfigure option. Result: The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks.

(Drivers 6.21g and later) Discovering LUNs

There are two ways to discover new LUNs and validate that they are visible on the host with driver versions 6.21g and later. You use one of the following methods:

- ◆ Reboot the host with the reconfigure option:
touch /reconfigure/sbin/init 6
- ◆ Dynamically discover LUNs by running `hbacmd rescanluns` on each target/initiator pair.

To use `hbacmd` to discover the new LUNs, complete the following steps.

Step	Action
1	Run <code>hbacmd</code> to get the list of initiators. Enter: hbacmd list Result: The software displays the list of adapters. Record the WWPNS of each target so that you can use it in the next step.
2	Run <code>hbacmd</code> to get the list of targets attached to the given initiator. Enter: hbacmd allnodelist <initiator_wwpn> Result: The software displays the list of targets currently bound to the given initiator. Record the WWPNS of each target so that you can use it in the next step.
3	Run <code>hbacmd</code> to rescan the LUNs attached to the given initiator/target pair. Enter: hbacmd rescanluns <initiator_wwpn> <target_wwpn> Result: The target/initiator pair you specify are rescanned. All the discovered LUNs are reported to the SCSI driver.
4	Repeat steps 2 and 3 for each target/initiator combination.

(Sun native drivers) Configuring the host to discover LUNs

Getting the controller number

Before you discover the LUNs, you need to determine what the controller number of the HBA is. To do this, complete the following step.

1	Use the <code>cfgadm -al</code> command to determine what the controller number of the HBA is. If you use the <code>/usr/sbin/cfgadm -c configure cx</code> command to discover the LUNs, you need to replace <code>x</code> with the HBA controller number. Example: To determine what the controller number of the HBA is, use the <code>cfgadm -al</code> command. To make this information easier to read, the key lines in this example are in bold. <pre>\$ cfgadm -al Ap_Id Type Receptacle Occupant Condition c0 fc-fabric connected configured unknown c0::500a098187f93622 disk connected configured unknown c0::500a098197f93622 disk connected configured unknown c1 scsi-bus connected configured unknown c1::dsk/c1t0d0 disk connected configured unknown c1::dsk/c1t1d0 disk connected configured unknown c2 fc-fabric connected configured unknown c2::500a098287f93622 disk connected configured unknown c2::500a098297f93622 disk connected configured unknown</pre>
---	--

Discovering LUNs

To have the host discover new LUNs and validate that they are visible on the host, complete the following step.

Step	Action
1	To discover new LUNs, enter <code>/usr/sbin/cfgadm -c configure cx</code> where <code>x</code> is the controller number of the HBA where the LUN is expected to be visible. If you do not see the HBA in the output, check your driver installation to make sure it is correct. Result: The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks.

Labeling the LUN and configuring Volume Management software

Labeling the new LUN

The format utility is a menu-driven script that is provided on the Solaris host. Use the format utility to format and label new LUNs. Follow the step below to activate the format utility and configure LUNs.

Step	Action
1	To label LUNs for the Solaris host: 1. <code>/usr/sbin/format</code> 2. At the <code>format></code> prompt, select the disk you want to modify. 3. When prompted to label the disk, enter y. The LUN is now labeled and ready for the volume manager to use. 4. When you finish, you can use the quit option to exit the utility. The following examples show the type of output you would see on a system using LPFC drivers and on a system using Sun native drivers.

Step	Action
	Example 1: This example labels disk number 5 on a system using LPFC drivers. (Portions of this example have been removed to make it easier to review.) <pre># format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c3t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@0,0 1. c4t0d0 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128> /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,0 2. c4t0d1 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128> /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,1 3. c4t0d2 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128> /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,2 4. c4t0d3 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128> /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,3 Specify disk (enter its number): 1 selecting c4t0d0 [disk formatted] ... Disk not labeled. Label it now? y</pre>

Step	Action
	Example 2: This example labels disk number 2 on a system that uses Sun native drivers. (Portions of this example have been removed to make it easier to review.) <pre> \$ format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e01008eb71,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e0100c6631,0 2. c6t500A098387193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098387193622,0 3. c6t500A098197193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098197193622,0 4. c6t500A098187193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098187193622,0 5. c6t500A098397193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd 16 sec 2048> /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098397193622,0 Specify disk (enter its number): 2 selecting c6t500A098387193622d0: TESTER [disk formatted] ... Disk not labeled. Label it now? y </pre>

Configuring volume management

When your configuration uses volume management software, you must configure the LUNs so they are under the control of the volume manager.

- ◆ Use the Veritas `vxdctl enable` command to bring new LUNs under Veritas control
- ◆ Use the Veritas `vxdiskadm` utility to manage existing disk groups

For additional information, refer to the documentation that shipped with your volume management software.

About this chapter

This chapter describes how to enable a host using the MPxIO stack to discover LUNs on the storage system.

Note

For general information on working with LUNs, see “[Basic Tasks for Working with LUNs](#)” on page 91.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Configuring the host to discover LUNs](#)” on page 110
- ◆ “[Completing the configuration](#)” on page 111

Configuring the host to discover LUNs

Getting the controller number

Before you discover the LUNs, you need to determine what the controller number of the HBA is. To do this, complete the following step.

1

Use the `cfgadm -al` command to determine what the controller number of the HBA is. If you use the `/usr/sbin/cfgadm -c configure cx` command to discover the LUNs, you need to replace x with the HBA controller number.

Example: To determine what the controller number of the HBA is, use the `cfgadm -al` command.

```
$ cfgadm -al
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::500a098187f93622	disk	connected	configured	unknown
c0::500a098197f93622	disk	connected	configured	unknown
c1	scsi-bus	connected	configured	unknown
c1::dsk/c1t0d0	disk	connected	configured	unknown
c1::dsk/c1t1d0	disk	connected	configured	unknown
c2	fc-fabric	connected	configured	unknown
c2::500a098287f93622	disk	connected	configured	unknown
c2::500a098297f93622	disk	connected	configured	unknown

Discovering LUNs

To have the host discover new LUNs and validate that they are visible on the host, complete the following step.

Step	Action
1	To discover new LUNs, enter the <code>/usr/sbin/cfgadm -c configure cx</code> command where <code>x</code> is the controller number of the HBA where the LUN is expected to be visible. If you do not see the HBA in the output, check your driver installation to make sure it is correct. Result: The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks.

Completing the configuration

**Formatting or
labeling the new
LUN**

The format utility is a menu-driven script that is provided on the Solaris host. Use the format utility to format and label LUNs.

Follow the step below to activate the format utility and configure LUNs.

Step	Action
1	To activate the format utility, enter the following command line from the host console: 1. <code>/usr/sbin/format</code> 2. Follow screen directions to format or label LUNs. Result: The format script displays available disks and configuration options. The following examples show the type of output you would see on a system using SPARC processors and on a system using x86/64 (AMD) processors.

Step	Action
	Example 1: This example labels disk number 15 on a SPARC system. (Portions of this example have been removed to make it easier to review.) <pre> \$ format Searching for disks...done AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107> /pci@1f,700000/scsi@2/sd@0,0 1. c1t1d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107> /pci@1f,700000/scsi@2/sd@1,0 2. c1t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107> flasharc /pci@1f,700000/scsi@2/sd@2,0 3. c1t3d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107> flasharc /pci@1f,700000/scsi@2/sd@3,0 4. c5t60A9800043346859444A2D36702D6B69d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d36702d6b69 5. c5t60A9800043346859444A2D36702D6D77d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d36702d6d77 ... 10. c5t60A9800043346859444A2D36702D787Ad0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128 /scsi_vhci/ssd@g60a9800043346859444a2d36702d787a 11. c5t60A9800043346859444A2D36702D6956d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d36702d6956 12. c5t60A9800043346859444A2D3670414A6Bd0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d3670414a6b 13. c5t60A9800043346859444A2D3670414D30d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d3670414d30 14. c5t60A9800043346859444A2D3670414F48d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d3670414f48 15. c5t60A9800043346859444A2D3670415A70d0 <NETAPP-LUN-0.2 cyl 2046 alt 2 hd 16 sec 128> /scsi_vhci/ssd@g60a9800043346859444a2d3670415a70 Specify disk (enter its number): 15 selecting c5t60A9800043346859444A2D3670415A70d0 [disk formatted] ... Disk not labeled. Label it now? y </pre>

Step	Action
	Example 2: This example runs the fdisk command and then labels disk number 15 on an x86/x64 (AMD) system. You must run the fdisk command before you can label a LUN. <pre> \$Specify disk (enter its number): 15 selecting C4t60A9800043346859444A2D367047492Fd0 [disk formatted] FORMAT MENU: disk - select a disk type - select (define) a disk type partition - select (define) a partition table current - describe the current disk format - format and analyze the disk fdisk - run the fdisk program repair - repair a defective sector label - write label to the disk analyze - surface analysis defect - defect list management backup - search for backup labels verify - read and display labels save - save new disk/partition definitions inquiry - show vendor, product and revision volname - set 8-character volume name !<cmd> - execute <cmd>, then return quit format> label Please run fdisk first. format> fdisk No fdisk table exists. The default partition for the disk is: a 100% "SOLARIS System" partition Type "y" to accept the default partition, otherwise type "n" to edit the partition table. y format> label Ready to label disk, continue? y format> </pre>

Configuring volume management

If your configuration uses either the Solaris Volume Manager, the ZFS file system, or the Veritas Volume Manager, you must configure the LUNs so they are under the volume manager control. See the documentation for the your volume manager for information on doing this.

Note

To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.

About this chapter

This chapter describes how to use the `sanlun` utility to display information about the LUNs, paths to the LUNs, and the host HBAs.

The `sanlun` utility is a tool provided by the SAN Toolkit that helps collect and report information about paths to your devices and how they map to LUNs on the storage system. You can also use this command to display information about the host HBAs.

Topics in this chapter

Topics in this chapter include the following:

- ◆ [“Displaying host LUN information”](#) on page 116
- ◆ [“Displaying information about the storage system paths”](#) on page 119
- ◆ [“Displaying host HBA information”](#) on page 124

Displaying host LUN information

Displaying host LUN information

To use `sanlun` to display information about the LUNs connected to the host, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
3	Enter the <code>sanlun lun show</code> command to display LUN information. The command has the following format: <pre>sanlun lun show [-v] [-d host device filename all storagesystem name storagesystem name:storagesystem pathname] -p</pre> <code>-v</code> produces verbose output. <code>-d</code> is the device option and can be one of the following: ❖ <code>host device filename</code> is the special device file name for the disk on Solaris (which might represent a storage system LUN). ❖ <code>all</code> lists all storage system LUNs under <code>/dev/rdsd</code>. ❖ <code>storage system name</code> lists all storage system LUNs under <code>/dev/rdsd</code> on that storage system. ❖ <code>storage system name:storage system pathname</code> lists all storage system LUNs under <code>/dev/rdsd</code> that are connected to the storage system path name LUN on that storage system. <code>-p</code> provides information about the primary and secondary paths available to the LUN when you are using multipathing. You cannot use the <code>-d</code> option if you use <code>-p</code>. Use the following format: <pre>sanlun lun show -p [storagesystem name:storagesystem pathname storagesystem name all]</pre> If you enter <code>sanlun lun show</code>, <code>sanlun lun show -p</code>, or <code>sanlun lun show -v</code> without any parameters, the utility responds as if you had included the <code>all</code> parameter. Result: The requested LUN information is displayed.

Step	Action
	For example, you might enter: ◆ sanlun lun show -p to display a listing of all the paths associated with the LUN. This information is useful if you need to set up path ordering or troubleshoot a problem with path ordering. ◆ sanlun lun show -d /dev/rdisk/<x> to display the summary listing of the LUN(s) associated with the host device /dev/rdisk<x> where <i>x</i> is a device such as /dev/rdisk/c#t#d#. ◆ sanlun lun show -v all to display verbose output for all the LUN(s) currently available on the host. ◆ sanlun lun show toaster to display a summary listing of all the LUNs available to the host served by the storage system called toaster. ◆ sanlun lun show toaster:/vol/vol0/lun0 to display a summary listing of all the LUNs available to the host served by lun0 on toaster. Note_____ When you specify either the <code>sanlun lun show <storage_system_name></code> or the <code>sanlun lun show <storage_system_name:storage_system_pathname></code> command, the utility displays only the LUNs that have been discovered by the host. LUNs that have not been discovered by the host are not displayed. _____

Step	Action
	Example: The following is an example of the output you see when you use the <code>sanlun show</code> command in verbose mode on a system using the Veritas DMP stack. <pre># ./sanlun lun show -v filer: lun-pathname device filename adapter lun size lun state filerX: /vol/vol1/hostA_lun2 /dev/rdsk/c0t500A098487093F9Dd5s2 emlxs0 3g (3221225472) GOOD Serial number: HnTMWZEqHyD5 Filer FCP nodename:500a098087093f9d Filer FCP portname:00a098397093f9d Filer adapter name: 0a Filer IP address: 10.60.181.66 Filer volume name:vol1 FSID:0x331ee81 Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 LUN partition table permits multiprotocol access: no why: bad starting cylinder or bad size for data partition LUN has valid label: yes</pre>

Displaying information about the storage system paths

Displaying path information

To display information about paths to the storage system, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the /opt/NTAP/SANToolkit/bin directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
3	At the host command line, enter the following command to display LUN information: sanlun lun show -p all -v -p provides information about the optimized (primary) and non-optimized (secondary) paths available to the LUN when you are using multipathing. Note _____ (MPxIO stack) MPxIO makes the underlying paths transparent to the user. It only exposes a consolidated device such as /dev/rdisk/c7t60A980004334686568343655496C7931d0s2. This is the name generated using the LUN's serial number in the IEEE registered extended format, type 6. The Solaris host receives this information from the SCSI Inquiry response. As a result, sanlun cannot display the underlying multiple paths. Instead it displays the target port group information. You can use the mpathadm or luxadm command to display the information if you need it. See “(MPxIO) Check that MPxIO and ALUA are working on the host” on page 135. _____ all lists all storage system LUNs under /dev/rdisk. -v turns on verbose mode. Because MPxIO provides multiple paths, sanlun selects one of the available paths to display. This is not a fixed path and could vary each time you execute the sanlun command with the -v option.

Step	Action
	Example: The following example uses the <code>-p</code> option to display information about the paths. <pre># sanlun lun show -p ONTAP_PATH: fas3020-rtp21:/vol/vol0/lun0 LUN: 0 LUN Size: 20g (21474836480) Host Device: /dev/rdisk/c7t60A980004334686568343655496C7931d0s2 LUN State: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: Native Multipath-provider: Sun Microsystems TPGS flag: 0x10Filer Status: TARGET PORT GROUP SUPPORT ENABLED Target Port Group : 0x1001 Target Port Group State: Active/optimized Vendor unique Identifier : 0x10 (2GB FC) Target Port Count: 0x2 Target Port ID : 0x1 Target Port ID : 0x2 Target Port Group : 0x3002 Target Port Group State: Active/non-optimized Vendor unique Identifier : 0x30 (2GB FC)</pre>

Step	Action
	Target Port Count: 0x2 Target Port ID : 0x101 Target Port ID : 0x102 Target Port Group : 0x3002 Target Port Group State: Active/non-optimized Vendor unique Identifier : 0x30 (2GB FC) Target Port Count: 0x2 Target Port ID : 0x101 Target Port ID : 0x102 ONTAP_PATH: fas3020-rtp21:/vol/vol1/lun1 LUN: 1 LUN Size: 20g (21474836480) Host Device: /dev/rdisk/c7t60A980004334686568343655496D5464d0s2 LUN State: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: Native Multipath-provider: Sun Microsystems TPGS flag: 0x10Filer Status: TARGET PORT GROUP SUPPORT ENABLED Target Port Group : 0x1001 Target Port Group State: Active/optimized Vendor unique Identifier : 0x10 (2GB FC) Target Port Count: 0x2 Target Port ID : 0x1 Target Port ID : 0x2 Target Port Group : 0x3002 Target Port Group State: Active/non-optimized Vendor unique Identifier : 0x30 (2GB FC) Target Port Count: 0x2 Target Port ID : 0x101 Target Port ID : 0x102

Interpreting `sanlun lun show -p` output

The `sanlun lun show -p` command provides the following information:

- ◆ **(MPxIO stack)** The target port group and its state, count, and ID. See information on ALUA in the *Data ONTAP® Block Access Management Guide for FCP*.
- ◆ **(MPxIO stack)** The vendor unique identifier
- ◆ **(Veritas DMP stack)** path state—Whether the path is enabled or disabled.
- ◆ **(Veritas DMP stack)** path type:.
 - ❖ Primary paths communicate directly using the adapter on the local storage system.
 - ❖ Secondary paths are proxied to the partner storage system over the cluster interconnect.
 - ❖ Standby occurs when the path is being serviced by a partner storage system in takeover mode. Note that this case occurs only when Veritas assumes the array policy is Active/Active. If the array policy is Active/Passive and the path is being served by the partner file in the takeover mode, that path state displays as secondary.
- ◆ device filename—The special device file name for the disk on Solaris that represents the LUN.
- ◆ host HBA—The name of the initiator HBA on the host.
- ◆ local storage system port—The port that provides direct access to a LUN. This port appears as the Xa for F8xx, gF8xx, FAS9xx, gFAS9xx, FAS3000, or FAS6000 storage systems and X_C for FAS270 storage systems, where X is the slot number of the HBA. This is always a primary path.
- ◆ partner storage system port—The port that provides passive path failover. This port appears as Xb for F8xx or FAS9xx storage systems and X_C for FAS270 on storage systems, where X is the slot number of the HBA. This is always a secondary path.

After the failover of a storage system cluster, the `sanlun lun show -p` command reports secondary paths as secondary but enabled, because these are now the active paths.

- ◆ Multipathing policy—The multipathing policy is one of the following:
 - ❖ **(Veritas DMP stack) A/A (Active/Active)**—Veritas DMP uses more than one path concurrently for traffic. The A/A policy indicates that the B ports on the storage system HBAs are in standby mode and become active only during a takeover. This means that there are two active paths to the cluster at any given time.
 - ❖ **(Veritas DMP stack) A/P-C (Active/Passive Concurrent)**—Groups of LUNs connected to a single controller failover as a single failover entity. Failover occurs at the controller level and not at the LUN level.

- ❖ **(MPxIO stack) Native**—this is displayed as the multipathing policy when MPxIO with ALUA is used for Multipathing.

Displaying host HBA information

Displaying host HBA information

To manage the host HBA information, complete the following steps.

Step	Action
1	Ensure that you are logged in as root on the host.
2	Change to the /opt/NTAP/SANToolkit/bin directory. Example: You can use the cd command to change to the:/opt/NTAP/SANToolkit/bin directory. # cd /opt/NTAP/SANToolkit/bin

Step	Action
3	At the host command line, enter the following command to display host HBA information: <pre>./sanlun fcp show adapter [-c] [-v] [adapter name all]</pre> -c option produces configuration instructions. -v option produces verbose output. all lists information for all FC adapters. adapter name lists information for a specified adapter. Result: The FC adapter information is displayed. Example: The following command line displays information about the adapter on a system using the qlc driver. <pre># ./sanlun fcp show adapter -v</pre> <pre>adapter name: qlc0 WWPN: 21000003ba16dec7 WWNN: 20000003ba16dec7 driver name: 20060630-2.16 model: 2200 model description: 2200 serial number: Not Available hardware version: Not Available driver version: 20060630-2.16 firmware version: 2.1.144 Number of ports: 1 port type: Private Loop port state: Operational supported speed: 1 GBit/sec negotiated speed: 1 GBit/sec OS device name: /dev/cfg/c1</pre>

Step	Action
	adapter name: qlc1 WWPN: 210000e08b88b838 WWNN: 200000e08b88b838 driver name: 20060630-2.16 model: QLA2462 model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel serial number: Not Available hardware version: Not Available driver version: 20060630-2.16 firmware version: 4.0.23 Number of ports: 1 of 2 port type: Fabric port state: Operational supported speed: 1 GBit/sec, 2 GBit/sec, 4 GBit/sec negotiated speed: 4 GBit/sec OS device name: /dev/cfg/c2 adapter name: qlc2 WWPN: 210100e08ba8b838 WWNN: 200100e08ba8b838 driver name: 20060630-2.16 model: QLA2462 model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel serial number: Not Available hardware version: Not Available driver version: 20060630-2.16 firmware version: 4.0.23 Number of ports: 2 of 2 port type: Fabric port state: Operational supported speed: 1 GBit/sec, 2 GBit/sec, 4 GBit/sec negotiated speed: 4 GBit/sec OS device name: /dev/cfg/c3

About this chapter

This appendix provides tips and troubleshooting information to help you as you use the FC Solaris Host Utilities. This appendix contains the following:

- ◆ Suggestions for checking your system.
- ◆ Information about possible problems and how to handle them.
- ◆ Diagnostic tools that you can use to gather information about your system.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Confirming that your system is set up correctly”](#) on page 128
- ◆ [“Using Diagnostic Utilities”](#) on page 140
- ◆ [“Collecting storage system information”](#) on page 141
- ◆ [“Collecting Solaris information”](#) on page 143
- ◆ [“Collecting switch information”](#) on page 144
- ◆ [“Known issues with the Host Utilities”](#) on page 146

Confirming that your system is set up correctly

Check the Release Notes for the most up-to-date information

Check the *Release Notes* for a list of known problems and limitations. The *Release Notes* also contain information on how to look up information about known bugs.

You can access the *Release Notes* from the NOW site. A link to the *Release Notes* is online in the SAN/IPSAN library at <http://now.netapp.com/NOW/knowledge/docs/san/>. A link is also available on the NOW Description Page for the Solaris Host 5.0 Utilities.

About the sections that follow

If you have any problems with the Host Utilities, you should make sure your system setup is correct.

The information in the sections that follow helps you verify your system setup. As you go through it, keep the following in mind:

For more information on any of the Solaris commands, see the man pages and operating system documentation.

For more information on the Data ONTAP commands, see the Data ONTAP documentation, in particular, the FCP and iSCSI Block Access Guide.

- ◆ You perform some of these checks from the host and others from the storage system. In some cases, you must have the Host Utilities SAN Toolkit installed before you perform the check. For example, the SAN Toolkit contains both the `basic_config` command and the `sanlun` command.
- ◆ To make sure you have the current version of the system components, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/l>. Support for new components is added on an on-going basis. This online document contains a complete list of supported HBAs, platforms, applications, and drivers. These pages are updated on a regular basis to ensure that they always have information on the most current components.

Note

The examples that follow were accurate at the time this document was produced. They are not intended to show all the options available.

Remove earlier versions of the software

If you currently have the FC Solaris Host Utilities 4.x or FC Solaris Attach Kit installed, you must remove it. See “[Removing an FC Solaris Host Utilities or Attach Kit installation](#)” on page 49.

(MPxIO) Ensure MPxIO is configured correctly

Configurations that use native drivers with MPxIO in a NetApp clustered environment require that ALUA be enabled. If your system was used for iSCSI or was part of a single NetApp storage controller configuration, the `symmetric-option` may have been set. This option disables ALUA on the host. To enable ALUA, you must remove this option by either:

- ◆ Commenting out the appropriate section in the `/kernel/drv/scsi_vhci.conf` file. The example below displays the section you must comment out.
- ◆ Running the `mpxio_set -d` command to automatically comment out the section.

After you comment out `symmetric-option`, you must reboot your system for the change to take effect.

Example: The following is the section of the `/kernel/drv/scsi_vhci.conf` file that you must comment out in order to enable MPxIO to work with ALUA.

```
#device-type-scsi-options-list =  
#"NETAPP LUN", "symmetric-option";  
#symmetric-option = 0x1000000;
```

Check the version of your host operating system

Make sure you have the correct version of the operating system. Use the `cat /etc/release` command to display information about your operating system.

Examples: This example checks the version on a SPARC system.

```
# cat /etc/release  
  
Solaris 10 5/08 s10s_u5wos_10 SPARC  
Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.  
Use is subject to license terms.  
Assembled 24 March 2008
```

This example checks the operating system version on an x86/x64 (AMD) system:

```
# cat /etc/release  
  
Solaris 10 5/08 s10x_u5wos_10 X86  
Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.  
Use is subject to license terms.  
Assembled 24 March 2008
```

Check the HBA

Supported HBAs should be installed before you install the Host Utilities. If you want to confirm that you have the supported HBAs, install the Host Utilities SAN Toolkit software package. Then use its `sanlun` command to display information about HBAs on your system.

If you are using a MPxIO stack, you can also use the `fcinfo hba-port` command to get information about the HBA.

Example 1: The following is an example of using the `sanlun` command to check a QLogic HBA using a Sun native qlc driver.

```
# sanlun fcp show adapter -v
```

```
adapter name:      qlc1
WWPN:              210000e08b88b838
WWNN:              200000e08b88b838
driver name:       20060630-2.16
model:             QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:     Not Available
hardware version:  Not Available
driver version:    20060630-2.16
firmware version:  4.0.23
Number of ports:   1 of 2
port type:         Fabric
port state:        Operational
supported speed:   1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed:  4 GBit/sec
OS device name:    /dev/cfg/c2
```

```
adapter name:      qlc2
WWPN:              210100e08ba8b838
WWNN:              200100e08ba8b838
driver name:       20060630-2.16
model:             QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:     Not Available
hardware version:  Not Available
driver version:    20060630-2.16
firmware version:  4.0.23
Number of ports:   2 of 2
port type:         Fabric
port state:        Operational
supported speed:   1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed:  4 GBit/sec
OS device name:    /dev/cfg/c3
```

Example 2: The following is an example using the `san lun` command to check your HBA using the `san lun` command when you have LPFC drivers.

```
# sanlun fcp show adapter -v
```

```
adapter name:      lpfc0
WWPN:              10000000c9512706
WWNN:              20000000c9512706
driver name:       lpfc
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM55099793
hardware version:  1036406d
driver version:    6.21g; HBAAPI(I) v2.0.f, 01-23-06
firmware version:  2.70A5 (B2F2.70A5)
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec
negotiated speed:  2 GBit/sec
OS device name:    /devices/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1

adapter name:      lpfc1
WWPN:              10000000c9512707
WWNN:              20000000c9512707
driver name:       lpfc
model:             LP11002-M4
model description: Emulex LP11002-M4 4Gb 2port FC: PCI-X2 SFF HBA
serial number:     VM55099793
hardware version:  1036406d
driver version:    6.21g; HBAAPI(I) v2.0.f, 01-23-06
firmware version:  2.70A5 (B2F2.70A5)
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec
negotiated speed:  2 GBit/sec
OS device name:    /devices/pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1,1
```

(MPxIO) Ensure that MPxIO is enabled on SPARC systems

When you use a MPxIO stack on a SPARC FC system, you must manually enable MPxIO. Verify whether MPxIO is enabled. If it is not, then you must manually enable using the `stmsboot` command. You must reboot your system after you run the `stmsboot` command.

The options you use with the `stmsboot` command vary depending on your version of Solaris:

- ◆ Solaris 10 update 5, execute the `stmsboot -D fp -e` command.
- ◆ Prior to Solaris 10 update 5, execute the `stmsboot -e` command.

You must reboot your system after you execute the `stmsboot` command.

Note

On x86/64 (AMD) systems, MPxIO is enabled by default.

Example: If MPxIO is not enabled on a system running Solaris 10 update 5, enter the following commands to enable MPxIO and then reboot your system.

```
# stmsboot -D fp -e
# touch /reconfigure
# init 6
```

Result: The `stmsboot -D fp -e` command enables MPxIO by setting the `mpxio_disable=no` option in the `fp.conf` file. It updates `/etc/vfstab`. The commands that follow cause the system to reboot.

(MPxIO) Check the Solaris Volume Manager

The MPxIO version of the Host Utilities supports the Solaris Volume Manager (SVM). To quickly check the condition of SVM volumes, enter the following command:

```
# metastat -a
```

In a Sun Cluster environment, metaset and their volumes are only displayed on the node that is controlling the storage.

See your Solaris documentation for more information on working with the SVM.

Check the Veritas Volume Manager

The Host Utilities supports the Veritas Volume Manager (VxVM) for the Veritas DMP stack and certain configurations of the MPxIO stack.

Note

(MPxIO) To determine which versions of VxVM are supported with MPxIO, see the NetApp Interoperability Matrix, at <http://now.netapp.com/matrix/mtx/login.do/>.

To quickly check the VxVM disks, enter the `vxdisk list` command:

Example: The following is an example of using the `vxdisk list` command. This example has been truncated to make it easier to read.

```
# vxdisk list
DEVICE      TYPE      DISK      GROUP      STATUS
Disk_0      auto:none  -         -         online invalid
Disk_1      auto:none  -         -         online invalid
FAS30200_0  auto:cdsdisk disk0     dg         online
FAS30200_1  auto:cdsdisk disk1     dg         online
FAS30200_2  auto:cdsdisk disk2     dg         online
FAS30200_3  auto:cdsdisk disk3     dg         online
...
#
```

To view volume information, use the `vxprint` command.

Example: The following is an example of using the `vxprint` command. This example has been truncated to make it easier to read.

```
# vxprint -g scdg
TY NAME          ASSOC      KSTATE  LENGTH  PLOFFS  STATE   TUTILO  PUTILO
dg scdg          scdg        -        -        -        -        -        -

dm disk00        Disk_49    -        10475520 -        -        -        -
dm disk01        Disk_33    -        10475520 -        -        -        -
dm disk02        Disk_51    -        10475520 -        -        -        -
dm disk03        Disk_54    -        10475520 -        -        -        -
dm disk04        Disk_17    -        10475520 -        -        -        -
dm disk05        Disk_39    -        10475520 -        -        -        -
dm disk06        Disk_58    -        10475520 -        -        -        -
dm disk07        Disk_24    -        10475520 -        -        -        -
dm disk08        Disk_11    -        10475520 -        -        -        -
dm disk09        Disk_34    -        10475520 -        -        -        -
dm disk10        Disk_35    -        10475520 -        -        -        -
dm disk11        Disk_29    -        10475520 -        -        -        -
dm disk12        Disk_53    -        10475520 -        -        -        -
dm disk13        Disk_2     -        10475520 -        -        -        -
dm disk14        Disk_12    -        10475520 -        -        -        -
dm disk15        Disk_5     -        10475520 -        -        -        -
dm disk16        Disk_44    -        10475520 -        -        -        -
dm disk17        Disk_32    -        10475520 -        -        -        -
dm disk18        Disk_10    -        10475520 -        -        -        -
dm disk19        Disk_27    -        10475520 -        -        -        -
v  raw_vol       fsgen      ENABLED  209510400 -        ACTIVE  -        -
pl raw_vol-01    raw_vol    ENABLED  209510400 -        ACTIVE  -        -
sd disk00-01     raw_vol-01 ENABLED  10475520 0        -        -        -
sd disk05-01     raw_vol-01 ENABLED  10475520 10475520 -        -        -
sd disk10-01     raw_vol-01 ENABLED  10475520 20951040 -        -        -
sd disk15-01     raw_vol-01 ENABLED  10475520 31426560 -        -        -
sd disk01-01     raw_vol-01 ENABLED  10475520 0        -        -        -
sd disk06-01     raw_vol-01 ENABLED  10475520 10475520 -        -        -
sd disk11-01     raw_vol-01 ENABLED  10475520 20951040 -        -        -
sd disk16-01     raw_vol-01 ENABLED  10475520 31426560 -        -        -
sd disk02-01     raw_vol-01 ENABLED  10475520 0        -        -        -
sd disk07-01     raw_vol-01 ENABLED  10475520 10475520 -        -        -
sd disk12-01     raw_vol-01 ENABLED  10475520 20951040 -        -        -
sd disk17-01     raw_vol-01 ENABLED  10475520 31426560 -        -        -
sd disk03-01     raw_vol-01 ENABLED  10475520 0        -        -        -
sd disk08-01     raw_vol-01 ENABLED  10475520 10475520 -        -        -
sd disk13-01     raw_vol-01 ENABLED  10475520 20951040 -        -        -
sd disk18-01     raw_vol-01 ENABLED  10475520 31426560 -        -        -
```

See your Veritas documentation for more information on working with the VxVM.

Check settings in `ssd.conf` or `sd.conf`

Verify that you have the correct settings in the `ssd.conf` file or in the `sd.conf` file.

On a SPARC system using MPxIO, running the `basic_config -ssd_set` command modifies the `/kernel/drv/ssd.conf` file.

On an x86/x64 (AMD) system using MPxIO, running the `basic_config -sd_set` command modifies the `sd.conf` file.

If you have a Veritas DMP stack, see “[Veritas DMP: Configuring the Solaris Host](#)” on page 53 for information on the required system parameters.

If you have a MPxIO stack, see “[MPxIO: Configuring the Solaris Host](#)” on page 65 for information on the required system parameters.

Example of `ssd.conf` file: (MPxIO on a SPARC system) You can confirm the `ssd.conf` file was correctly set up by checking it to ensure that it contains the following:

```
ssd-config-list="NETAPP  LUN", "netapp-ssd-config";
netapp-ssd-
config=1,0x9007,64,300,30,0,0,0,0,0,0,0,0,0,0,0,30,0,0,8,0,0;
```

Example of sd.conf file: (MPxIO on an x86/64 system) You can confirm the ss.conf file was correctly set up by checking it to ensure that it contains the following:

```
sd-config-list="NETAPP LUN", "netapp-sd-config";
netapp-sd-
config=1,0x9c01,64,0,0,0,0,0,0,0,0,0,0,300,30,30,0,0,8,0,0;
```

(MPxIO) Check that MPxIO and ALUA are working on the host

Confirm that you have MPxIO and ALUA working on the host. You can use either the `luxadm display` command or the `mpathadm` command to check for multipathing and the `sanlun lun show -p` command to check for ALUA.

Note.

You must enable ALUA before you map the LUNs. If you do not enable ALUA, MPxIO will not work. For more information on using ALUA, see the *Data ONTAP® Block Access Management Guide for FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. In particular, see the section “*Enabling ALUA for a Fibre Channel igroup.*”

Example of checking for MPxIO using luxadm: If MPxIO and ALUA are set up correctly, you should see some primary and secondary paths when you run the `luxadm display` command and the state of these paths should be listed as **online**. (To make this information easier to locate in the following example, these lines are shown in bold.)

```
# luxadm display /dev/rdisk/c2t60A98000686F646E4B5A34784A477668d0s2
DEVICE PROPERTIES for disk:
/dev/rdisk/c2t60A98000686F646E4B5A34784A477668d0s2
Vendor:                NETAPP
Product ID:            LUN
Revision:              0.2
Serial Num:            hodnKZ4xJGvh
Unformatted capacity: 1024.000 MBytes
Read Cache:            Enabled
  Minimum prefetch:    0x0
  Maximum prefetch:    0x0
Device Type:           Disk device
Path(s):
/dev/rdisk/c2t60A98000686F646E4B5A34784A477668d0s2
/devices/scsi_vhci/ssd@g60a98000686f646e4b5a34784a477668:c,raw
Controller             /devices/pci@1d,700000/fibre-
channel@2,1/fp@0,0
  Device Address          500a098185447022,9
  Host controller port WWN 10000000c94be343
  Class                 secondary
  State                 ONLINE
Controller              /devices/pci@1d,700000/fibre-channel@2/fp@0,0
  Device Address          500a098195447022,9
  Host controller port WWN 10000000c94be342
  Class                 primary
  State                 ONLINE
```

Example of checking for MPxIO using mpathadm: You can use the `mpathadm list lu` command to get a list of all the LUNs.

```
# mpathadm list lu
/dev/rdisk/c3t60A980004334612F466F4C6B72483362d0s2
  Total Path Count: 8
  Operational Path Count: 8
/dev/rdisk/c3t60A980004334612F466F4C6B72483230d0s2
  Total Path Count: 8
  Operational Path Count: 8
/dev/rdisk/c3t60A980004334612F466F4C6B7248304Dd0s2
  Total Path Count: 8
  Operational Path Count: 8
/dev/rdisk/c3t60A980004334612F466F4C6B7247796Cd0s2
```

```

Total Path Count: 8
Operational Path Count: 8
/dev/rdisk/c3t60A980004334612F466F4C6B72477838d0s2
Total Path Count: 8
Operational Path Count: 8
/dev/rdisk/c3t60A980004334612F466F4C6B72477657d0s2
Total Path Count: 8
Operational Path Count: 8

```

Example of checking for ALUA on the host: If ALUA is correctly set up, this information should appear when you run the `sanlun lun show -p` command.

```

# sanlun lun show -p
ONTAP_PATH:shf820-01:/vol/vol1/san_sol1
LUN:0
LUN Size: 20g (21474836480)
Host Device:/dev/rdisk/c4t600A09804F6443496434345435586476d0s2
lun state: GOOD Filer_CF_State: Cluster Enabled
Multipath_Policy: Native Multipath-provider: Sun Microsystems
Filer Status: TARGET PORT GROUP SUPPORT ENABLED
Target Port Group : 0x1
Target Port Group State: Active/optimized
Vendor unique Identifier : 0x0 (Other FC)
Target Port Count: 0x4
    Target Port ID : 0x1
    Target Port ID : 0x2
    Target Port ID : 0x3
    Target Port ID : 0x4
Target Port Group : 0x2
Target Port Group State: Active/non-optimized
Vendor unique Identifier : 0x0 (Other FC)
Target Port Count: 0x4
    Target Port ID : 0x101
    Target Port ID : 0x102
    Target Port ID : 0x103
    Target Port ID : 0x104

```

Check the storage system setup

You need to have your storage system set up correctly. For information on what needs to be done, see the list of storage system requirements in the section [“Prerequisites”](#) on page 11.

**(Veritas DMP)
Check that the ASL
and APM have been
installed**

You must have the ASL and APM installed in order for the Veritas DMP to identify whether the path is primary or secondary. Without the ASL and APM, the DSM treats all paths as equal, even if they are secondary paths. As a result, You may see I/O errors on the host. On the storage system, you may see the Data ONTAP error: FCP Partner Path Misconfigured. If you encounter these errors, make sure you have the ASL and APM installed.

Check the NetApp Interoperability Matrix Tool at
<http://now.netapp.com/matrix/mtx/>

**(MPxIO) Check the
ALUA settings on
the storage
system**

You must have ALUA set on the storage system to work with igroups. To make sure you have ALUA set for your igroup, use the `igroup show -v` command.

Example: The following command line displays information about the `cfmode` on the storage system and shows that ALUA is enabled. (To make the information on ALUA easier to locate, it is shown in bold.)

```
filerA# igroup show -v
shsun07 (FCP):
    OS Type: solaris
    Member: 10:00:00:00:c9:4b:e3:42 (logged in on: 0c)
    Member: 10:00:00:00:c9:4b:e3:43 (logged in on: vtic)
    ALUA: Yes
```

Note

For more information on ALUA, see the *Data ONTAP® Block Access Management Guide for FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. In particular, see the section “*Enabling ALUA for a Fibre Channel igroup.*”

**Verify that the
switch is installed
and configured**

If you have a fabric-attached configuration, check that the switch is set up and configured as outlined in the Installation and Setup instructions that shipped with your hardware. You should have completed the following tasks:

- ◆ Installed the switch in a system cabinet or rack.
- ◆ Checked that the switch firmware is supported with this Host Utility. See the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.
- ◆ Turned on power to the switch.
- ◆ Configured the network parameters for the switch, including its serial parameters and IP address.

Determine whether to use switch zoning

If you have a fabric-attached configuration, determine whether switch zoning is appropriate for your system setup.

Zoning requires more configuration on the switch, but it results in the following:

- ◆ Configuring the host is simplified.
- ◆ Information becomes more manageable. The output from the host tool `iostat` is easier to read because fewer paths are displayed.

To have a high-availability configuration, make sure that each LUN has at least one primary path and one secondary path through each switch. For example, if you have two switches, you would have a minimum of four paths per LUN.

It is recommended that your configuration have no more than eight paths per LUN. For more information about zoning, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.

For more information on supported switch topologies, see the *iSCSI and Fibre Channel Configuration Guide*, available on the NOW site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Power up equipment in the correct order

The different pieces of hardware communicate with each other in a prescribed order, which means that problems occur if you turn on power to the equipment in the wrong order. Throughout the procedures that follow, this is the order you must use when powering on the equipment:

- ◆ Configured Fibre Channel switches
It can take several minutes for the switches to boot.
- ◆ Disk shelves
- ◆ Storage systems
- ◆ Host

Verify that the host and storage system can communicate

Once your setup is complete, make sure the host and the storage system can communicate.

You can verify that the host can communicate with the storage system by issuing a command from:

- ◆ Storage system's console
- ◆ A remote login that you established from the host

Using Diagnostic Utilities

Diagnostic utilities

The SAN Toolkit includes several command-line utilities that provide diagnostic information you can use for troubleshooting.

When you run these utilities, they gather configuration information about different aspects of your system. They automatically collect information and place it in a directory. Then they create a compressed tar file of that directory. You can send this file to NetApp customer support.

Normally, you only use these utilities when requested to do so by Customer Support. The utilities are:

- ◆ `filer_info`. See “[Collecting storage system information](#)” on page 141 for more information.
- ◆ `solaris_info`. See “[Collecting Solaris information](#)” on page 143 for more information.
- ◆ `<switch>_info`, where `<switch>` is the brand of switch you are using (for example, `brocade_info`). See “[Collecting switch information](#)” on page 144 for more information.

General steps for executing the programs

Although some of the scripts have arguments that you may want to use, the general steps for executing these scripts are the same.

Step	Action
1	Change to the <code>opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
2	Enter the utility name. If the utility you are using includes arguments, you can enter them as well. Result: The configuration information that the utility gathers is placed in a compressed file.

Collecting storage system information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run `filer_info` to collect storage system configuration information. You should only run this utility when requested to do so by Customer Support. To execute `filer_info`, complete the following steps.

Collecting storage system diagnostics

Complete the steps below to collect storage system information.

Step	Action
1	On the Solaris host, change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
2	At the host command prompt, enter the following command: <pre>filer_info [-d <i>path_and_file_name</i>] [-l <i>passwd</i>] -n <<i>storage_system_name</i>> -f</pre> -d enables you to specify a different path in which to place the resulting output. -l <i>passwd</i> allows you to specify a password to access the storage system. -n <<i>storage_system_name</i>> specifies the target of the command, where <i>storage_system_name</i> is the name of IP address of the storage system accessed. -f automatically overwrites an existing <<i>storage_system_name</i>> file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the <i>filer_info</i> script tells you to delete or rename the file. Result: The output of various status commands is placed in the output directory, where <i>storage_system_name</i> is replaced by the storage system name or IP address. The contents of this directory are placed in a tarred and compressed file named <i>storage_system_<storage_system_name>.tar.Z</i>.
3	Send the resulting compressed file to Customer Support for analysis.

Collecting Solaris information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run `solaris_info` to collect configuration information about your host. You should only run this utility when requested to do so by Customer Support.

Collecting Solaris diagnostics

To execute `solaris_info`, complete the following steps:

Step	Action
1	On the Solaris host, change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>
2	At the command prompt, enter the following command: <code>solaris_info [-d <directory>] [-n <file_name>] -f</code> -d <directory> specifies the location for the file that <code>solaris_info</code> creates. -n <file_name> specifies the name for the output file and directory names. -f automatically overwrites an existing <code>.tar.Z</code> file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the script tells you to delete or rename the file. Result: When you run <code>solaris_info</code> , the utility gathers information about your Solaris configuration and places it in a <code>.tar.Z</code> file. It also creates a directory and populates it with files containing output from numerous commands run on the Solaris system and contents of the <code>/var/adm/message</code> file.
3	Send the resulting file to Customer Support for analysis.

Collecting switch information

When to run the diagnostic script

If you are having a problem, Customer Support might ask you to run the switch utility script that is provided when you install the SAN Toolkit software. A utility is provided for each switch that is supported for the FC Solaris Host Utilities. The utility name is <switch>_info, where *switch* is the brand of the switch (for example, if you have a McDATA switch, the file name would be mcddata_info). This utility collects switch configuration information. You should only run this utility when requested to do so by Customer Support.

Collecting switch diagnostics

To execute the utility, complete the following steps.

Step	Action
1	Change to the /opt/NTAP/SANToolkit/bin directory: cd /opt/NTAP/SANToolkit/bin

Step	Action
2	At the command prompt, enter the following command: <pre><switch_brand>_info [-d <i>path_and_file_name</i>] [-l <i>user:password</i>] -n <switchname> -f</pre> <switch_brand> is the type of the switch you are using; for example, brocade for a Brocade switch. The -d <i>path_and_file_name</i> option enables you to specify a different path in which to place the resulting output. The -l <i>user:password</i> option allows the choice of a user name and password on the switch. If this option is not specified, the Brocade default settings of admin:password are used for Brocade switches. The target of the command is specified by -n <i>switchname</i>, where <i>switchname</i> is the name or IP address of the switch to be interrogated. -f automatically overwrites an existing output file in the default output directory. You might have an existing file in this directory if you previously ran the script. If you do not use the -f option, and you have an existing file, the script tells you to delete or rename the file. Result: The output of various status commands on the switch is placed in a default directory where <i>switch_brand</i> is the brand of the switch, such as Brocade, and <i>switchname</i> is replaced by the switch name or IP address. The contents of this directory are placed in a file named <switch_brand>_switch_switchname.tar.Z. Example: This example uses the brocade_info tool. The switch brcd1 is connected with the user “john”, who has the password “smith”. <pre>host1> brocade_info -d /opt/NTAP/SANToolkit/bin/ -l john:smith -n brcd1</pre> The information from the switch is placed in the brocade_switch_brcd1.tar.Z file.
3	Send the resulting file to Customer Support for analysis.

Known issues with the Host Utilities

(MPxIO) sanlun doesn't show all adapters

Symptom: The `sanlun lun show` command does not display all the adapters.

```

hostA# sanlun lun show
filer:                lun-pathname                device filename
adapter protocol      lun size          lun state
fas3020-rtp22: /vol/vol14/lun45
/dev/rds/c7t60A9800043346859444A36554A4D6E53d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol13/lun44
/dev/rds/c7t60A9800043346859444A36554A4D5765d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol11/lun42
/dev/rds/c7t60A9800043346859444A36554A4C775Ad0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol10/lun41
/dev/rds/c7t60A9800043346859444A36554A4C642Fd0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol9/lun40
/dev/rds/c7t60A9800043346859444A36554A4C4B75d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol8/lun39
/dev/rds/c7t60A9800043346859444A36554A4B7175d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol7/lun38
/dev/rds/c7t60A9800043346859444A36554A4B6243d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp2: /vol/vol6/lun37
/dev/rds/c7t60A9800043346859444A36554A4B486Ed0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol4/lun35
/dev/rds/c7t60A9800043346859444A36554A4A684Bd0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol2/lun33
/dev/rds/c7t60A9800043346859444A36554A4A3736d0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol1/lun32
/dev/rds/c7t60A9800043346859444A36554A496B6Bd0s2 qlc1      FCP      20g
(21474836480)      GOOD
fas3020-rtp22: /vol/vol0/lun31

```

Explanation: When you use MPxIO, there are multiple paths. MPxIO controls the path over which the sanlun SCSI commands are sent and it uses the first one it finds. This means that the adapter name can vary each time you issue the `sanlun lun show` command.

What to do: To display information on all the adapters, use either the `luxadm display` command or the `mpathadm` command. For the `luxadm display` command, you would enter:

```
luxadm display -v device_name
```

Where *device_name* is the name of the device you are checking.

(MPxIO) Solaris log message says data not standards compliant

Symptom: The Solaris log contains a message stating:

WARNING: Page83 data not standards compliant

Explanation: This Solaris log message sometimes occurs when you are executing the Host Utilities. It is Solaris bug and has been reported to Sun. The Solaris initiator implements an older version of the SCSI Spec.

What to do: The NetApp SCSI target is standards compliant, so ignore this message.

Appendix overview

This appendix describes how to use the `create_binding.pl` script to create persistent bindings. This script is bundled with the FC Solaris Host Utilities software.

Note

You only need to create persistent binding if you are using a host running the Veritas DMP stack with LPFC. If you are using the MPxIO stack, you do not need to create persistent bindings.

You can also create bindings using `hbacmd` or HBAware.

- ◆ To create persistent bindings with `hbacmd`, refer to “[Creating WWPN bindings with hbacmd](#)” on page 101.
- ◆ To create persistent bindings with HBAware, refer to HBAware documentation

Topics in this appendix

This appendix includes the following topics:

- ◆ “[The create_binding.pl script](#)” on page 150
- ◆ “[Before you begin](#)” on page 152
- ◆ “[Using the create_binding.pl script](#)” on page 154

The create_binding.pl script

What is the create_binding.pl script

The `create_binding.pl` script is a NetApp-provided tool that enables you to bind the WWPNs of a single storage system's active ports to a target ID. The tool then saves these bindings in the `/kernel/drv/lpfc.conf` file.

When you upgrade an installation that uses WWNN bindings to WWPN bindings, the `create_binding.pl` script saves the existing `/kernel/drv/lpfc.conf` file to `/kernel/drv/lpfc.conf.NTAP.[date-stamp]`. The script then deletes the WWNN or WWPN bindings in the `/kernel/drv/lpfc.conf` file and replaces them with WWPN bindings to the new target IDs. The previous WWNN or WWPN bindings remain in the `/kernel/drv/lpfc.conf` file but are commented out. You can refer to the commented-out settings for information about your previous bindings.

For new installations in which there are no existing WWPN bindings, the `create_binding.pl` script accesses the storage system's WWPN list and automatically binds the WWPNs to target IDs and updates the `/kernel/drv/lpfc.conf` file.

Rules the script uses for WWPN bindings

The script uses the following rules for creating WWPN bindings:

- ◆ Initiators bind to accessible targets only.
- ◆ For multipath configurations with Veritas VxVM and DMP software, the script automatically creates bindings for all accessible paths to the storage system.
- ◆ On non-multipathing hosts, an initiator cannot bind to more than one active port per storage system. For these configurations, the script prompts you to select one binding for one accessible path to the storage system for each HBA.
- ◆ When the storage system's FC cluster failover mode (`cfmode`) is set to standby, the script creates bindings for only the active A ports on the storage systems in a cluster. Standby mode means that the B ports are not available to the initiator until a takeover occurs.
- ◆ When the storage system's FC cluster failover mode is set to single image, the script creates bindings for all available target ports.
- ◆ The script has two options for configurations that have multiple paths to the storage system.
 - d—enables you to configure bindings for all available paths when Veritas VxVM with DMP is installed.

-D—enables you to configure bindings for all available paths when Veritas VxVM with DMP is not yet installed.

Caution

All multipath configurations require Veritas VxVM with DMP. If you use the -D option to create bindings and you do not have Veritas software, you must install Veritas software after you create bindings. You cannot create or use LUN data until Veritas software is installed.

Before you begin

Prerequisites for using the `create_binding.pl` script

To use the `create_binding.pl` script, your configuration must meet the following prerequisites:

- ◆ The LPFC driver is installed.
- ◆ The FC Solaris Host Utilities is installed.

The `create_binding.pl` script is bundled with the FC Solaris Host Utilities software. You must install this kit to obtain the `create_binding.pl` software.

Refer to “[Installing the Host Utilities software package](#)” on page 42 for information about downloading and installing the FC Solaris Host Utilities.

- ◆ Cabling is connected.

The `create_binding.pl` script checks the availability of paths from the host to the storage system. All cables must be connected correctly between the host system and the storage system.

- ◆ Zoning is configured.

If zoning is configured on your switches, the zones must be configured correctly and must be active. You can verify that zones are configured correctly in the following ways:

- ❖ By checking the zoning information on the switch to verify that the storage system ports and the host ports are visible to the switch.
- ❖ By verifying that storage system ports are visible to the host. You can run the `fcplib show initiator` command on the storage system console to view the WWPNs of initiators connected to the storage system’s target HBAs.

- ◆ Storage clusters are configured.

The `create_binding.pl` script configures WWPN bindings for both storage systems in a cluster configuration. If you have a storage cluster configuration, the cluster must be set up and configured before you run the script.

Storage system prerequisites

To use the `create_binding.pl` script, your storage system must meet the following prerequisites:

- ◆ Remote Shell (`rsh`) service is installed.
The script uses the `rsh` service to access the storage system. You enable the `rsh` service by entering the following command:
`options rsh.enable.on`
- ◆ The FC protocol service is licensed and started on the storage system.
- ◆ The SAN Solaris host is a trusted host of the storage system. You use the following storage system command to make the SAN Solaris host a trusted host:
`options trusted.hosts <hostname>`

Using the create_binding.pl script

Creating WWPN bindings with create_binding.pl

To use the `create_binding.pl` script create WWPN bindings, complete the following steps.

Step	Action
1	Ensure that you are logged in to the host with root privileges.
2	Change to the <code>/opt/NTAP/SANToolkit/bin</code> directory: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
3	Enter the following command: <pre>create_binding.pl -l 'login:passwd,login:passwd' -n filer [-d -D]</pre> -l <i>login:passwd</i> is the login name and password used to log on to the storage system. Use a quoted, comma separated pair to specify two different login:passwd strings for a storage system and its partner. -n <i>storage_system</i> is the storage system's host name. For cluster configurations, you specify one member of the cluster. The script detects and configures bindings for the partner storage system. -d enables the script to process information for Veritas DMP configurations. Use the -d option if you have already installed Veritas Volume Manager (VxVM) with DMP. If you use the -d option for Veritas DMP configurations, the script automatically configures WWPN bindings for all accessible paths to the storage system. -D enables the script to create bindings for all accessible paths to the storage system when Veritas DMP has not yet been installed. Use this option if you have multiple paths to the storage system, but have not installed Veritas VxVM with DMP. Caution The -D option enables you to configure persistent bindings even if you have not yet installed Veritas software. If you have multiple paths to the storage system, you must install Veritas software after you configure persistent bindings. Multipathing software is required for configurations with multiple paths to the storage system. Example : <pre>./create_binding.pl -l root -n storage_system_1 -d</pre>

Step	Action
4	Take one of the actions below: ◆ If you have a single-attached or direct-attached configuration, the script displays the new WWPN bindings. • Enter y to save the new WWPN bindings and proceed to Step 5 ◆ If you have VxVM with DMP and you did not use the -d option to the <code>create_binding.pl</code> script in Step 3, the script displays a warning because you did not use the -d option. The <code>create_binding.pl</code> script will not discover multiple paths to the storage system: • Enter n to exit the script • Rerun the <code>create_binding.pl</code> script with the -d option • Enter y at the prompt to continue the script and bind by WWPN Result: The script displays the new WWPN bindings and prompts you to save them to the <code>/kernel/drv/lpfc.conf</code> file. • Enter y to save the new WWPN bindings • Proceed to Step 5 ◆ If you have Veritas VxVM and DMP installed without multipathing configuration and you did not use the -d option of the script in Step 3, the script displays a warning message. • At the warning, enter y to continue. • For each HBA, choose one of the available paths to the storage system. Each path is a WWPN binding. • When prompted, enter y to save the new WWPN bindings to the <code>/kernel/drv/lpfc.conf</code> file. • Proceed to Step 5.

Step	Action
5	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for the target ID mapped to the storage system's WWPN. The entry has the following format: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> You must create an entry for every LUN mapped to the host. Note_____ If the host is connected to multiple storage systems or clusters, you must create entries in sd.conf for each additional storage system or cluster, and for each LUN mapped to the host. _____
6	On the host console, reboot the host with the reconfigure option: <pre>touch /reconfigure</pre> <pre>init 6</pre> Note_____ If you're performing multiple configuration steps that require rebooting, you may want to perform all those steps and then reboot. _____

About this appendix This appendix describes what you need to do to configure SAN boot on a Solaris host running the Veritas DMP stack.

Note

NetApp qualifies solutions and components on an ongoing basis. To verify that SAN booting is supported in your configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do>.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[Configuration overview](#)” on page 160
- ◆ “[\(Sun native drivers\) Setting up the HBA](#)” on page 162
- ◆ “[Creating the bootable LUN](#)” on page 172
- ◆ “[Performing a Direct Installation to a SAN Boot LUN](#)” on page 178
- ◆ “[\(LPFC, Sun Native Drivers\) Copy bootblks, data to SAN boot LUN](#)” on page 179
- ◆ “[Modifying OpenBoot](#)” on page 183

Configuration overview

Contents of this appendix

This appendix provides steps for configuring a SAN boot LUN and installing a new OS image on it. You can apply these steps to most configurations.

Keep in mind that you can use other methods to create a SAN boot LUN. This appendix does not describe other examples of bootable LUNs, such as configurations that boot multiple volumes or use diskless servers.

If you are using Sun native drivers, refer to Solaris documentation for details about additional configuration methods. In particular, see Sun's document *Sun StorEdge SAN Foundation Software Configuration Guide*, which is available in PDF format on the Sun site.

Prerequisites

The procedures and examples in this appendix assume the following:

- ◆ The FC Solaris Host Utilities software has been installed and the host and storage system are configured properly and use software and firmware that is supported by the Host Utilities.
- ◆ The host operating system is installed on a local disk and uses a UFS file system.
- ◆ Bootcode/FCode is downloaded and installed on the HBA.
 - ❖ For Emulex HBAs, the FCode is available on the Emulex site. For details on how to load the FCode onto the HBA, follow the steps in [“Updating your FCode”](#) on page 36.
 - ❖ For QLogic-branded HBAs, the FCode is available on the QLogic site. For details on how to load the FCode onto the HBA, follow the steps in [“Upgrading the QLogic FCode”](#) on page 38.
 - ❖ For Sun-branded QLogic HBAs, the FCode is available as a patch from Sun. Follow steps in the patch Readme for details on how to load this FCode onto the HBA.
- ◆ If you are using Emulex HBAs, you must have the Emulex FCA utilities with EMLXemlxu and the EMLXdrv have been installed. For more information on these tools, see [“Installing the EMLXemlxu utilities”](#) on page 34.
- ◆ For Emulex-branded HBAs and Sun-branded Emulex HBAs, the FCode is available on the Emulex site. For details on how to load the FCode onto the HBA, follow the steps in [“Updating your FCode”](#) on page 36.

- ◆ If you are using QLogic-branded HBAs, you must have the SANsurfer SCLI utility installed. For more information on this tool, see “[Installing the SANsurfer CLI package](#)” on page 37.
- ◆ If you are using the Emulex LPFC driver, make sure you have configured persistent binding. There are three basic ways to configure the bindings:
 - ❖ Manually
 - ❖ By using `create_binding.pl`
 - ❖ By using `hbacmd`.

For information about host, storage system and firmware versions, refer to the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Configuration overview

Configuring a bootable LUN involves:

1. Making sure the HBA is set to the appropriate Bootcode

For more information on setting this, see the chapter “[Veritas DMP: Managing the Driver Software](#)” on page 19.
2. Creating the boot LUN
 - ❖ Create the LUN that you will use for the bootable LUN.
 - ❖ Display the size and layout of the partitions of the current Solaris boot drive
 - ❖ Partition the bootable LUN to match the host boot drive
3. Selecting the method for installing to the SAN booted LUN
 - ❖ Directly Install the bootblks onto the bootable LUN
 - ❖ Perform a file system dump and restore to a SAN booted LUN
 1. Install the bootblks onto the bootable LUN
 2. Copy the boot data from the source disk onto the bootable LUN
4. Modifying Bootcode
 - ❖ Verify the OpenBoot version
 - ❖ Set the FC topology to the bootable LUN
 - ❖ Bind the adapter target and the bootable LUN
 - ❖ Create an alias for the bootable LUN
5. Rebooting the system.

(Sun native drivers) Setting up the HBA

Setting up the HBA

Part of configuring a bootable LUN involves setting up your HBA. You may need to shut down the system and switch the HBA mode.

Note

The steps in this section apply only to the SUN native HBAs. If you are using LPFC HBAs, see the section “(LPFC drivers) Setting up the HBA with Emulex drivers” on page 168.

Changing the HBA mode

The steps you need to perform depend on the type of HBA you are using. For information on changing the mode on:

- ◆ Emulex HBAs, see “Changing the Emulex HBA to SFS mode” on page 162
- ◆ QLogic HBAs, see “Changing the QLogic HBA to enable FCode compatibility” on page 166

Changing the Emulex HBA to SFS mode

To change the mode on an Emulex HBA from an SD compatible mode to an SFS mode, you must bring the system down and then change each HBA. Complete the following steps.

WARNING

These steps will change the device definition from lpfc@ to emlxs@. Doing this will cause the controller instance to be incremented. Any devices current on the controllers that are being modified will receive new controller numbers. If you are currently mounting these devices via the /etc/vfstab file, these entries will be invalid.

Step	Action
1	At the operating system prompt, issue the <code>init 0</code> command. <code># init 0</code>
2	When the ok prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>

Step	Action
4	Issue the <code>show-devs</code> command to see the current device names. Example: Executing the <code>show-devs</code> shows that the Emulex device has not been set to SFS mode. If the mode had been changed, the devices would have displayed <code>.../emlxs</code> instead of <code>.../lpfc</code>. See Step 5 for information setting the devices to SFS mode. <pre data-bbox="194 352 662 937">ok show-devs controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/lpfc@3 /pci@8,700000/lpfc@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3</pre>

Step	Action
5	Select the first Emulex device and set it to SFS mode using the <code>set -sfs-boot</code> command. Doing this changes the devices to emlx devices. Example: The first command, <code>show-devs</code>, shows the device name. The next command, <code>select</code>, selects the device <code>lpfc@0</code>. The command, <code>set -sfs-boot</code>, sets the SFS mode. Some of the output has been truncated to make this easier to read. To make the SFS information easy to locate, it is in bold. <pre> ok show-devs ...output truncated... /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0,1 /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0 ...output truncated... ok select /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0 ok set-sfs-boot Flash data structure updated. Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 1 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 0 WWPN 0000.0000.0000.0000 LUN (8 bytes) 0000.0000.0000.0000 *** Type reset-all to update. *** ok </pre>
6	Repeat the previous step for each Emulex device.
7	Enter the <code>reset-all</code> command to update the devices. Example: The <code>reset-all</code> command updates the Emulex devices with the new mode. <pre> ok reset-all Resetting ... </pre>

Step	Action
8	Issue the show-devs command to confirm that you have changed the mode on all the Emulex devices. Example: Executing the show-devs command lets you confirm that the Emulex devices are showing up as emlx devices. To continue the example in Step 5, the device selected there, /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0, has been changed to an emlx device. In a production environment, you would want to change all the devices to emlx. <pre> ok show-devs /pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0,1 /pci@0/pci@0/pci@8/pci@0/pci@1/emlx@0 /memory-controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>
9	Set the auto-boot? back to true and boot the system with a reconfiguration boot. Example: The boot command brings the system back up. <pre> ok setenv auto-boot? true ok boot - r </pre>

Changing the QLogic HBA to enable FCode compatibility

To enable FCode compatibility on a QLogic HBA, you must bring the system down and then change each HBA. Complete the following steps.

Step	Action
1	At the operating system prompt, issue the <code>init 0</code> command. # <code>init 0</code>
2	When the <code>ok</code> prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>
4	Issue the <code>show-devs</code> command to see the current device names. Example: Execute the <code>show-devs</code> shows to see whether there is FCode compatibility. <code>ok show-devs</code> ... <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0/disk</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0</code> <code>/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0/disk</code>
5	Select the first QLogic device. Example: Use the <code>select</code> command to select the first QLogic device. <code>ok select /pci@7c0/pci@0/pci@8/QLGC,qlc@0,1</code> QLogic QLE2462 Host Adapter Driver(SPARC): 1.16 03/10/06
6	Execute the <code>set-mode</code> command and, if needed, set the compatibility mode to <code>fcode</code> . Example: The <code>set-mode</code> command sets the mode to <code>fcode</code> . <code>ok set-mode</code> Current Compatibility Mode: <code>fcode</code> Do you want to change it? (y/n) Choose Compatibility Mode: <code>0 - fcode</code> <code>1 - bios</code> <code>enter: 0</code> Current Compatibility Mode: <code>fcode</code>

Step	Action
7	Execute the <code>set-fc-mode</code> command and, if needed, set the Fcode Mode to <code>qlc</code>. Example: The <code>set -mode</code> command sets the mode to <code>fcode</code>. <pre>ok set-fc-mode Current Fcode Mode: qlc Do you want to change it? (y/n) Choose Fcode Mode: 0 - qlc 1 - qla enter: 0 Current Fcode Mode: glc</pre>
8	Repeat the previous steps to configure each QLogic device.
9	Enter the <code>reset -all</code> command to update the devices. Example: The <code>reset -all</code> command updates the QLogic devices with the new mode. <pre>ok reset-all Resetting ...</pre>
10	Set the <code>auto-boot?</code> back to <code>true</code> and boot the system with a reconfiguration boot. Example: The <code>boot</code> command brings the system back up. <pre>ok setenv auto-boot? true ok boot -r</pre>

(LPFC drivers) Setting up the HBA with Emulex drivers

Setting up the LPFC HBA

Part of configuring a bootable LUN involves setting up your LPFC drivers. Before you do this, you need to shut down the system and set the Emulex HBA to SD mode.

Note — The steps in this section apply only to the LPFC HBAs. If you are using Sun native HBAs, see the section “Setting up the HBA (Sun native HBAs).”

WARNING — These steps may change the device definition from `lpfc emlxs@` to `lpfc@`. Doing this will cause the controller instance to be incremented. Any devices current on the controllers that are being modified will receive new controller numbers. If you are currently mounting these devices via the `/etc/vfstab` file, these entries will be invalid.

Changing the Emulex HBA to SD mode

To change the HBA from SFS compatible mode to SD mode, you must bring the system down and then change each HBA.

Complete the following steps.

Step	Action
1	At the operating system prompt, issue the <code>init 0</code> command. <code># init 0</code>
2	When the <code>ok</code> prompt appears, enter the <code>setenv auto-boot? false</code> command. <code>ok setenv auto-boot? false</code>
3	Enter the <code>reset-all</code> command. <code>ok reset-all</code>

Step	Action
4	Issue the <code>show-devs</code> command to see the current device names. Example: Executing the <code>show-devs</code> shows that the Emulex device has not been set to SD mode. If the mode had been changed, the devices would have displayed <code>.../lpfc</code> instead of <code>.../emlxs</code>. See Step 5 for information setting the devices to SD mode. <pre> ok show-devs controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@m0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/emlxs@3 /pci@8,700000/emlxs@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>

Step	Action
5	Select the first Emulex device and set it to SD mode using the <code>set -sd-boot</code> command. This changes the Emulex devices to lpfc devices. Example: The first command, <code>select</code>, selects the Emulex device 700000/emlxs@1. The second command, <code>set -sd-boot</code>, sets the SD mode. To make the SFS information easy to locate, it is in bold. <pre> ok select /pci@8,700000/lpfc@1 ok set-sd-boot Flash data structure updated. Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 0 WWPN 0000.0000.0000.0000 LUN (8 bytes) 0000.0000.0000.0000 *** Type reset-all to update. *** </pre>
6	Repeat the previous step for each Emulex device.
7	Enter the <code>reset -all</code> command to update the devices. Example: The <code>reset -all</code> command updates the Emulex devices with the new mode. <pre> ok reset-all Resetting ... </pre>

Step	Action
8	Issue the show-devs command to confirm that you have changed the mode on all the Emulex devices. Example: Executing the show-devs command lets you confirm that all the Emulex devices are showing up as lpfc devices. <pre> ok show-devs /pci@8,600000 /pci@8,700000 /memory-controller@1,400000 /SUNW,UltraSPARC-III+@1,0 /memory-controller@0,400000 /SUNW,UltraSPARC-III+@0,0 /virtual-memory /memory@0,0 /aliases /options /openprom /chosen /packages /pci@8,600000/SUNW,qlc@4 /pci@8,600000/SUNW,qlc@4/fp@0,0 /pci@8,600000/SUNW,qlc@4/fp@0,0/disk /pci@8,700000/lpfc@3 /pci@8,700000/lpfc@1 /pci@8,700000/scsi@6,1 /pci@8,700000/scsi@6 /pci@8,700000/usb@5,3 </pre>
9	Set the auto-boot? back to true and boot the system with a reconfiguration boot. Example: The boot command brings the system back up. <pre> ok setenv auto-boot? true ok boot -r </pre>

Creating the bootable LUN

Create the LUN

Use standard storage system commands and procedures to create and map a LUN assigned to a host.

For detailed information about configuring igroups and LUNs, refer to *Data ONTAP® Software Setup Guide* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Partition the bootable LUN to match the host device

You must partition the bootable LUN so that it matches the partitions on the host boot device. This involves:

- ◆ Displaying information about the host boot device
- ◆ Modifying the bootable LUN to model the partition layout of the host boot device

Complete the following steps to partition the bootable LUN.

Step	Action
1	Select the current Solaris boot device: - Execute the Solaris format utility. Enter /usr/sbin/format. Result: The system displays Available Disk Selections. Example: The format utility displays information about the disks on your system. <pre># format AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@0,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@1,0 2. c1t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@2,0 3. c1t3d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@3,0 4. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b 5. c2t500A098786F7CED2d12 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,c 6. c2t500A098786F7CED2d13 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,d 7. c2t500A098786F7CED2d14 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,e 8. c2t500A098786F7CED2d15 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,f</pre> Note_____ If you want to set up additional SAN boot configurations, you can use the <code>san lun</code> command to find the primary path. The examples in this section use the c1t1d0 device. _____ - At the specify disk prompt, enter the number that corresponds to the current Solaris boot device. This is the device that maps to the root file system in the <code>/etc/vfstab</code> file. Result: The system selects the Solaris boot disk and displays partition options

Step	Action
2	Look in the /etc/vfstab file and write down the information about the boot file systems. Example: The following example shows the contents of the /etc/vfstab file. Devices c1t1d0s0 and c1t1d0s1 (in bold) refer to the source boot device and must be replicated on the bootable LUN. <pre># vi /etc/vfstab # device device mount FS fsck mount mount # to mount to fsck point type pass at boot options # fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c1t1d0s1 - - swap - no - /dev/dsk/c1t1d0s0 /dev/rdisk/c1t1d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>

Step	Action																																																						
3	Display partition information for the current Solaris boot disk: - At the <code>format></code> prompt, enter p to display partition menu options - At the <code>partition></code> prompt, enter p to print the partition table for the current Solaris boot device. Write down the following information: ❖ Partition size: The bootable LUN must be the same size or larger than the original boot volume. If your system requires Veritas, you must allow an additional 1 GB for Veritas encapsulation. ❖ Partition layout: The bootable LUN must contain the same partitions in the same locations as the host boot device(s). The devices that you must include are referred to in the <code>/etc/vfstab</code> file. Result: The system displays the partitions for the boot device. Example: The following is the partition information for the current root disk slice being used for the SAN booting example in this chapter. Current partition table (original): Total disk cylinders available: 6142 + 2 (reserved cylinders) <table><tr><th>Part</th><th>Tag</th><th>Flag</th><th>Cylinders</th><th>Size</th><th>Blocks</th></tr><tr><td>0</td><td>root</td><td>wm</td><td>683 - 5461</td><td>14.00GB</td><td>(4779/0/0) 29362176</td></tr><tr><td>1</td><td>swap</td><td>wu</td><td>0 - 682</td><td>2.00GB</td><td>(683/0/0) 4196352</td></tr><tr><td>2</td><td>backup</td><td>wu</td><td>0 - 6141</td><td>17.99GB</td><td>(6142/0/0) 37736448</td></tr><tr><td>3</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>4</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>5</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>6</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr><tr><td>7</td><td>unassigned</td><td>wm</td><td>0</td><td>0</td><td>(0/0/0) 0</td></tr></table> - At the <code>partition></code> prompt, enter q to go back to the top level <code>format></code> prompt. Note _____ If your configuration boots off of more than one device, repeat steps 1 through 3 to determine the size and layout of each device. You must create and configure a bootable LUN that matches each boot device on the host. _____	Part	Tag	Flag	Cylinders	Size	Blocks	0	root	wm	683 - 5461	14.00GB	(4779/0/0) 29362176	1	swap	wu	0 - 682	2.00GB	(683/0/0) 4196352	2	backup	wu	0 - 6141	17.99GB	(6142/0/0) 37736448	3	unassigned	wm	0	0	(0/0/0) 0	4	unassigned	wm	0	0	(0/0/0) 0	5	unassigned	wm	0	0	(0/0/0) 0	6	unassigned	wm	0	0	(0/0/0) 0	7	unassigned	wm	0	0	(0/0/0) 0
Part	Tag	Flag	Cylinders	Size	Blocks																																																		
0	root	wm	683 - 5461	14.00GB	(4779/0/0) 29362176																																																		
1	swap	wu	0 - 682	2.00GB	(683/0/0) 4196352																																																		
2	backup	wu	0 - 6141	17.99GB	(6142/0/0) 37736448																																																		
3	unassigned	wm	0	0	(0/0/0) 0																																																		
4	unassigned	wm	0	0	(0/0/0) 0																																																		
5	unassigned	wm	0	0	(0/0/0) 0																																																		
6	unassigned	wm	0	0	(0/0/0) 0																																																		
7	unassigned	wm	0	0	(0/0/0) 0																																																		

Step	Action
4	Use format to create the root same size or larger and swap slice for the SAN boot disk. 1. At the format> prompt, type disk to get a list of disks. 2. At the specify disk prompt, enter the number that corresponds to the LUN mapped to host controller and target0 Note_____ Make a note of the identifying information for the LUN you select. You will use this information in “Modifying OpenBoot” on page 183, when you bind the bootable LUN to the adapter and create a boot alias. Pay close attention to the “cXtYdZ” information. Example: This example selects c2t500A098786F7CED2d11 as the bootable LUN. <pre>format> disk AVAILABLE DISK SELECTIONS: 0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@0,0 1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@1,0 2. c1t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@2,0 3. c1t3d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424> /pci@1f,700000/scsi@2/sd@3,0 4. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b 5. c2t500A098786F7CED2d12 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,c 6. c2t500A098786F7CED2d13 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,d 7. c2t500A098786F7CED2d14 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,e 8. c2t500A098786F7CED2d15 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,f Specify disk (enter its number): Specify disk (enter its number): 4 selecting c2t500A098786F7CED2d11 [disk formatted]</pre>

Step	Action
5	Display partition information for the current bootable LUN: 1. At the format> prompt, enter p to display partition menu options 2. At the partition> prompt, enter p to print the partition table for the current Solaris boot device. Result: The system displays the partitions for the bootable LUN.
6	At the partition> prompt, enter the number of the slice you want to modify to match the host boot device. Repeat this step for each slice. Set any unused slice to 0 cylinders in length.
7	At the partition> prompt, enter label l to write the new label to the disk.
8	Exit the format utility. 1. At the partition> prompt, enter q. 2. At the format> prompt, enter q.

Performing a Direct Installation to a SAN Boot LUN

Installing directly onto a SAN boot LUN

If you are using Sun native drivers, you can perform a direct installation onto the SAN boot LUN. You can install the operating system using:

- ◆ A CD/DVD installation
- ◆ A jumpstart installation

During the installation, select the device you created in the steps in the section [“Creating the bootable LUN”](#) on page 172.

(LPFC, Sun Native Drivers) Copy bootblks, data to SAN boot LUN

Install the bootblk The bootblk contains startup information, required by the Solaris host, that you must install onto the raw bootable LUN. Installing the bootblk involves:

- ◆ Using the `uname -a` command to determine the directory in `/usr/platform` where the bootblk is located
- ◆ Installing the bootblk onto the bootable LUN.

Complete the following steps to install the bootblk onto the bootable LUN.

Step	Action
1	Determine the directory in <code>/usr/platform</code> where the bootblks are located. Enter: uname -a Example: The output from the example that follows indicates that bootblks should reside in <code>/usr/platform/sun4u</code> (bold). <pre>#uname -a SunOS shsun17 5.10 Generic_118833-33 sun4u sparc SUNW,Sun-Fire-280R</pre>
2	Write the bootblk. Enter: <pre>/usr/sbin/installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/<device_name></pre> Example: The following example installs bootblks in the <code>/usr/platform/`uname -i`</code> directory on <code>/dev/rdisk/c2t500A098786F7CED2d11s0</code>. <pre>#!/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/c2t500A098786F7CED2d11s0</pre>

Copy the boot data Copying the boot data to the bootable LUN involves

- ◆ Creating and mounting a file system on the bootable LUN. The file system must match the file system type on the host boot device. The example in this document uses a UFS file system.
- ◆ Moving boot data from the host device onto the bootable LUN
- ◆ Editing the `vfstab` file to reference the bootable LUN

Complete the following steps to copy the boot data onto the bootable LUN.

Step	Action
1	Create a file system on the bootable LUN that matches the file system in the vfstab file of the source device. Enter: newfs <device> Example: The following creates a USF file system at /dev/rdisk/c2t500A098786F7CED2d11s0. <pre>newfs on /dev/rdisk/c2t500A098786F7CED2d11s0 newfs: /dev/rdisk/c2t500A098786F7CED2d11s0 last mounted as / newfs: construct a new file system /dev/rdisk/c2t500A098786F7CED2d11s0: (y/n)? y Warning: 4096 sector(s) in last cylinder unallocated /dev/rdisk/c2t500A098786F7CED2d11s0: 20971520 sectors in 3414 cylinders of 48 tracks, 128 sectors 10240.0MB in 214 cyl groups (16 c/g, 48.00MB/g, 5824 i/g) super-block backups (for fsck -F ufs -o b=#) at: 32, 98464, 196896, 295328, 393760, 492192, 590624, 689056, 787488, 885920, 20055584, 20154016, 20252448, 20350880, 20449312, 20547744, 20646176, 20744608, 20843040, 20941472</pre> The next example shows the corresponding entry (in bold) for the file system in the vfstab file. <pre>#device device mount FS fsck mount mount #to mount to fsck pont type pass at boot options fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c2t500A098786F7CED2d11s1 - - swap - no - /dev/dsk/c2t500A098786F7CED2d11s0 /dev/rdisk/c2t500A098786F7CED2d11s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>
2	Mount the root directory of the bootable LUN. Enter: mount <filesystem_name> /<mountpoint> Example: The following mounts the file system on the bootable LUN under /mnt. <pre>#mount /dev/dsk/c2t500A098786F7CED2d11s0 /mnt</pre>

Step	Action
3	Create the required directory structure onto the bootable LUN and copy the boot data. You need to perform a dump and a restore from the current boot disk to the SAN disk. You can use the commands: <pre> ufsdump 0f - <source_boot_device> (cd/ <mntpoint_of_bootable_lun>; ufsrestore rf -) </pre> Example: The following brings the system down to single user mode and performs a dump of the current boot disk. It then restores it to the SAN boot disk. <pre> #init S #df -k Filesystem kbytes used avail capacity Mounted on /dev/dsk/c1t1d0s0 10327132 3882023 6341838 38% / # ufsdump 0f - /dev/dsk/c1t1d0s0 (cd /mnt; ufsrestore rf -) DUMP: Date of this level 0 dump: Wed 28 Mar 2007 02:22:29 PM EDT DUMP: Date of last level 0 dump: the epoch DUMP: Dumping /dev/rdisk/c1t1d0s0 (sunv440-shu05:/) to standard output. DUMP: Mapping (Pass I) [regular files] DUMP: Mapping (Pass II) [directories] DUMP: Writing 32 Kilobyte records DUMP: Estimated 8004176 blocks (3908.29MB). DUMP: Dumping (Pass III) [directories] DUMP: Dumping (Pass IV) [regular files] Warning: ./lost+found: File exists DUMP: 98.97% done, finished in 0:00 DUMP: 8004158 blocks (3908.28MB) on 1 volume at 6506 KB/sec DUMP: DUMP IS DONE </pre>

Step	Action
4	Edit the /etc/vfstab file on the SAN boot LUN to refer to the correct device. You should use a primary to the LUN. To determine the primary path, run <code>san lun lun show -p</code> refer to the bootable LUN. The <code>san lun</code> command comes with the Host Utilities. Example: The following shows the vfstab entry for the bootable LUN (bold). <pre>#vi /mnt/etc/vfstab #more /mnt/etc/vfstab #device device mount FS fsck mount mount #to mount to fsck point type pass at boot options # fd - /dev/fd fd - no - /proc - /proc proc - no - /dev/dsk/c2t500A098786F7CED2d11s1 - - swap - yes - /dev/dsk/c2t500A098786F7CED2d11s0 /dev/rdsk/c2t500A098786F7CED2d11s0 / ufs 1 yes - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>

Modifying OpenBoot

What is OpenBoot

OpenBoot is the firmware that the host uses to start up the system. OpenBoot firmware also includes the hardware level user interface that you use to configure the bootable LUN.

**(Sun native drivers)
Modifying
OpenBoot**

Modifying OpenBoot when you are using Sun native drivers involves creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@8,700000/emlxs@1/sd@3,0 in the example that follows) during subsequent boot operations.

To modify OpenBoot, complete the following steps.

Step	Action
1	Activate the OpenBoot environment by performing the following steps: 1. Halt the Solaris operating system by entering the command: init 0 Result: The Open Boot program displays the OpenBoot prompt. 2. Stop automatic reboot during OpenBoot configuration by entering the command: setenv auto-boot? false 3. Reset the host and activate your changes by entering: reset-all

Step	Action
2	Run the show-disks command to get the correct device path name for your Fibre Channel card. It is important to look for the device path associated with the LUN you used during the earlier format step. Due to the way the OBP sees hardware, the device may show up with either “disk” or “ssd” as the terminating characters and those characters may differ from what was seen in the format output. Example: This output shown here continues the previous example and uses the information that was written down during the format command. The LUN is “w500a098786f7ced2,b” and the device path is “/pci@1d,700000/emlx@2/fp@0,0/ssd.” <pre>4. c2t500A098786F7CED2d11 <NETAPP-LUN-0.2 cyl 1022 alt 2 hd 16 sec 128> /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b</pre> The following example simply shows the type of output you see when you execute the show-disks command. <pre>{1} ok show-disks a) /pci@1f,700000/scsi@2,1/disk b) /pci@1f,700000/scsi@2/disk c) /pci@1e,600000/ide@d/cdrom d) /pci@1e,600000/ide@d/disk e) /pci@1d,700000/emlx@2,1/fp@0,0/disk f) /pci@1d,700000/emlx@2/fp@0,0/disk q) NO SELECTION Enter Selection, q to quit:</pre>
3	Create an alias for the bootable LUN. Enter: nvalias <i><alias_name></i> <i>/<boot_LUN_pathname></i> Example: This example creates an alias called “sanbootdisk”. <pre>{0} ok nvalias sanbootdisk /pci@1d,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b {0} ok nvstore {0} ok devalias sanbootdisk /pci@8,700000/emlx@2/fp@0,0/ssd@w500a098786f7ced2,b</pre>

Step	Action
4	To allow the system to autoboot to the LUN, add the LUN's alias to your boot device list. Enter the following command: <pre>setenv boot-device <alias_name></pre> Note_____ Before adding the new alias, use the <code>printenv boot-device</code> command to see if there are already boot devices in the list. If aliases are in the list, you'll need to include them on the <code>setenv boot-device</code> command line after the bootable LUN's alias. That alias must always be the first argument. The format for the command would be: <pre>setenv boot-device <bootable_lun_alias_name> <existing_alias1> <existing_alias2> ...</pre> _____ Example: Suppose you've created the alias <code>sanbootdisk</code> for your bootable LUN. You use the check the boot device list and discover that there are already two boot devices in the list: <code>disk</code> and <code>net</code>. To add the alias <code>sanbootdisk</code>, you would enter the <code>setenv boot-device</code> command followed by the <code>sanbootdisk</code> and the other two aliases in the list. <pre>{0} ok setenv boot-device sanbootdisk disk net</pre>
5	Enable auto-boot. Enter: <pre>setenv auto-boot? true</pre>
6	Boot the LUN. Enter: <pre>boot <alias_name></pre> Example: The following command boots the alias for <code>/pci@1d,700000/emlx@2/fp@0,0/disk@w500a098786f7ced2,b</code> <pre>{0} ok boot sanbootdisk</pre>
7	Once you have booted the host on the newly created SAN LUN, you should enable the multipathing that is appropriate for your system. See the vendor documentation for your multipathing solution for more information.

Modify OpenBoot for the Emulex LPFC driver

Modifying OpenBoot for LPFC involves

- ◆ Running `sanlun` to document the information about your LUN.
- ◆ Mapping the HBA instance to a physical path.
- ◆ Binding the target and the LUN to the HBA.

- ◆ Creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@1d,700000/lpfc@2/sd@0,f5 in the example that follows) during subsequent boot operations.

The examples that follow use this LUN information:

```
c5t0d245 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 256>
      /pci@1d,700000/lpfc@2/sd@0,f5
```

To modify OpenBoot, complete the following steps.

Step	Action
1	Run the <code>san lun lun show -p</code> command and write down the information for the SAN boot LUN. Make sure you write down the correct device name so that it matches the name you used when you made changes to the <code>/etc/vfstab</code> on the SAN boot LUN. Note_____ The <code>san lun</code> command is in the <code>/opt/NTAP/SANToolkit/bin/</code> directory. Example: The <code>san lun lun show -p</code> command produces the following information. <pre># /opt/NTAP/SANToolkit/bin/san lun lun show -p fas3020-shu01:/vol/vol245/lun245 (LUN -1) 12g (12884901888) LUN State: GOOD Filer_CF_State: Cluster Enabled Filer Partner: fas3020-shu02 Multipath_Policy: None Multipath-provider: None ----- host filer path path device host primary partner state type filename HBA filer filer ----- up secondary /dev/rdisk/c6t3d245s2 lpfc1 0d up secondary /dev/rdisk/c6t2d245s2 lpfc1 0c up primary /dev/rdisk/c6t1d245s2 lpfc1 0d up primary /dev/rdisk/c6t0d245s2 lpfc1 0c up secondary /dev/rdisk/c5t3d245s2 lpfc0 0b up secondary /dev/rdisk/c5t2d245s2 lpfc0 0a up primary /dev/rdisk/c5t1d245s2 lpfc0 0b up primary /dev/rdisk/c5t0d245s2 lpfc0 0a</pre>

Step	Action
2	Get the device path for the lpfc instance to device path mapping from the /etc/path_to_inst file. Example: You can use the grep command to get this information. <pre># grep \"lpfc\" /etc/path_to_inst "/pci@1d,700000/lpfc@2" 0 "lpfc" "/pci@1d,700000/lpfc@2,1" 1 "lpfc"</pre>
3	Document the WWPN binding for your controllers. The persistent binding information can be found in the file /kernel/drv/lpfc.conf. Look for the variable named fcp-bind-WWPN. You should also look for the lpfc instance and target combination to match the path used in the /etc/vfstab file on the SAN boot LUN. Example: The binding for the sample configuration is provided below: <pre>fcp-bind-wwpn="500a098486c7ced2:lpfc0t0", "500a098386c7ced2:lpfc0t1", "500a098286c7ced2:lpfc1t0", "500a098186c7ced2:lpfc1t1", "500a098496c7ced2:lpfc0t2", "500a098396c7ced2:lpfc0t3", "500a098296c7ced2:lpfc1t2", "500a098196c7ced2:lpfc1t3";</pre>
4	Activate the OpenBoot environment by performing the following steps: 1. Halt the Solaris operating system by entering the command: init 0 Result: The Open Boot program displays the OpenBoot prompt. 2. Stop automatic reboot during OpenBoot configuration by entering the command: setenv auto-boot? false 3. Reset the host and activate your changes by entering: reset-all

Step	Action
5	Bind the HBA to the LUN and the target. This step uses all of the information gathered in steps 1 through 3. To determine the value of <i>target-in-hex</i>, you need to refer to the persistent binding entries you documented in step 3. The target is the value found in position Y of <i>wwpn</i>: <i>lpfcXtY</i>. Convert this number to hex and use it in the step below. <pre>select /device_path_to_card wwpn <wwpn-of-target> <lun-in-hex> <target-in-hex> set-boot-id reset-all</pre> Example: The following example show the type of output that occurs. <pre>ok select /pci@1d,700000/lpfc@2 ok wwpn 500a098486c7ced2 f5 0 set-boot-id Flash data structure updated. Signature 4e45504f Valid_flag 14e Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 0 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) f5 DID 0 WWPN 500a.0984.86c7.ced2 LUN (8 bytes) 0000.0000.0000.0000 *** Type reset-all to update. *** ok reset-all</pre>

Step	Action
6	Run the show-disks command to get the correct device path name for your Fibre Channel card. It is important to look for the device path associated with the LUN you used during the earlier format step. Due to the way the OBP sees hardware, the device may show up with either “disk” or “sd” as the terminating characters and those characters may differ from what was seen in the format output. Example: As can be seen in the output below, the device uses "sd" to represent the disk device. <pre>ok show-disks a) /pci@1f,700000/scsi@2,1/disk b) /pci@1f,700000/scsi@2/disk c) /pci@1e,600000/ide@d/cdrom d) /pci@1e,600000/ide@d/disk e) /pci@1d,700000/lpfc@2,1/sd f) /pci@1d,700000/lpfc@2/sd q) NO SELECTION Enter Selection, q to quit: Enter Selection, q to quit:</pre>
7	Create an alias for the bootable LUN. Enter: nvalias <i><alias_name></i> <i></boot_LUN_pathname></i> Example: This example creates an alias called “sanbootdisk”. <pre>0} ok nvalias sanbootdisk /pci@1d,700000/lpfc@2/sd@0,f5 {0} ok nvstore {0} ok devalias sanbootdisk /pci@1d,700000/lpfc@2/sd@0,f5</pre>

Step	Action
8	To allow the system to autoboot to the LUN, add the LUN's alias to your boot device list. Enter the following command: <pre>setenv boot-device <alias_name></pre> Note_____ Before adding the new alias, use the <code>printenv boot-device</code> command to see if there are already boot devices in the list. If aliases are in the list, you'll need to include them on the <code>setenv boot-device</code> command line after the bootable LUN's alias. That alias must always be the first argument. The format for the command would be: <pre>setenv boot-device <bootable_lun_alias_name> <existing_alias1> <existing_alias2> ...</pre> _____ Example: Suppose you've created the alias <code>sanbootdisk</code> for your bootable LUN. You use the check the boot device list and discover that there are already two boot devices in the list: <code>disk</code> and <code>net</code>. To add the alias <code>sanbootdisk</code>, you would enter the <code>setenv boot-device</code> command followed by the <code>sanbootdisk</code> and the other two aliases in the list. <pre>{0} ok setenv boot-device sanbootdisk disk net</pre>
9	Enable auto-boot. Enter: <pre>setenv auto-boot? true</pre>
10	Boot the LUN. Enter: <pre>boot <alias_name></pre> Example: The following command boots the alias for <code>/pci@1d,700000/lpfc@2/sd@0,f5</code> <pre>{0} ok boot sanbootdisk</pre>
11	Once you have booted the host on the newly created SAN LUN, you should enable the multipathing that is appropriate for your system. See the vendor documentation for your multipathing solution for more information.

About this chapter

This appendix describes what you need to do to configure SAN booting on an Solaris host that is using the MPxIO stack.

Note

NetApp qualifies solutions and components on an ongoing basis. To verify that SAN booting is supported in your configuration, see the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[Configuration overview](#)” on page 192
- ◆ “[Performing a direct installation](#)” on page 194
- ◆ “[File system dump, restore for a SPARC system without MPxIO](#)” on page 195
- ◆ “[File system dump, restore for an x86/x64 system with MPxIO](#)” on page 200
- ◆ “[File system dump, restore for an x86/x64 system without MPxIO](#)” on page 204

Configuration overview

Contents of this appendix

This appendix provides steps for configuring a SAN boot LUN and installing a new OS image on it. You can apply these steps to most configurations. The example in this appendix uses a Solaris host to configure a SAN boot LUN and install a new OS image onto it. The procedure described here is one method for creating a bootable LUN. You can apply this procedure to most configurations.

This appendix does not describe other examples of bootable LUNs, such as configurations that boot multiple volumes or use diskless servers.

Refer to Solaris documentation for details about additional configuration methods.

Sun documentation for this procedure

For complete details on the procedures that follow, see Sun's document *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide*, which is available in PDF format on the Sun site at docs.sun.com.

Make sure, though, that you use the FCode specified in the NetApp Interoperability Matrix at <http://now.netapp.com/matrix/mtx/login.do/>.

Prerequisites

The procedures and examples in this chapter assume the following:

- ◆ The FC Solaris Host Utilities software has been installed and the host and storage system are configured properly and use software and firmware that is supported by the Host Utilities.
- ◆ The host operating system is installed on a local disk and uses a UFS file system.

For information about host, storage system and firmware versions, refer to the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Special note about device names

To get the correct device name to allow for a successful boot, use the device path as reported by the OpenBoot PROM (OBP). Execute the `show-disks` command to display this name. This is because the OBP name is not always the same as the Solaris driver name.

Examples in this appendix

This appendix provides examples of copying data from an existing, locally booted server. The examples show how to perform this task using:

- ◆ A direct installation
- ◆ A file system dump and restore for a SPARC system with MPxIO disabled
- ◆ A file system dump and restore for an x86/x64 (AMD) system with MPxIO disabled
- ◆ A file system dump and restore for an x86/x64 (AMD) system with MPxIO enabled

Performing a direct installation

Options for setting up SAN booting

There are two ways to set up SAN booting:

- ◆ Copying the data from a locally booted device to a SAN device. Then reconfigure the host to boot to the SAN device. On x86/x64 (AMD) systems, using this method means you do not have to re-setup a currently configured host and its MPxIO settings.
- ◆ Installing the operating system directly on the SAN device. Once you do this, you must configure the SAN boot LUN on the storage system and map it to the host.

Performing a direct install

To perform a direct install, run the Solaris Installer and choose the SAN device as your install device.

File system dump, restore for a SPARC system without MPxIO

SPARC systems without MPxIO: Copying data from locally booted server

The following steps show you how to perform a file system dump and restore for a SPARC system that has MPxIO disabled. If you have a SPARC system and you want to copy data from a locally booted server, complete the following steps.

Step	Action
1	If MPxIO is enabled, disable it: ◆ For Solaris 10 update 5 and later, enter: <code># stmsboot -D fp -d</code> ◆ For Solaris 10 releases prior to update 5, enter: <code># stmsboot -d</code>
2	Use the mount command to identify the current boot device. Example: The following mounts the current boot device. <pre># mount / on /dev/dsk/c0t0d0s0 read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=8000 10 on Tue Jan 2 07:44:48 2007</pre>
3	Configure the LUN for ◆ The root file system ◆ The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices. In addition, you might want to set up the swap slice and slices for the metadata database replicas. Use the format command to choose the LUN. Write down the device path for the LUN so you can refer to it later. Example: In this example, you first run the format command. Based on its output, you determine that you want device 4, which has device path <code>/pci@1d,700000/emlx@2,1/fp@0,0/ssd@w500a09838350b481,f4</code>. <pre>4. c1t500a09838350b481d244 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 256> /pci@1d,700000/emlx@2,1/fp@0,0/ssd@w500a09838350b481,f4</pre>

Step	Action
4	Install the bootblock on the new bootable LUN by using the <code>installboot</code> command: <pre>installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/<boot lun></pre> Example: The following command line installs the bootlock on the boot LUN c1t500A09838350B481d32s0. <pre>installboot /usr/platform/'uname -i'/lib/fs/ufs/bootblk /dev/rdisk/c1t500A09838350B481d32s0</pre>
5	Create the SAN file system that you will dump the current file system to. Example: The following command line creates a new file system. <pre>newfs /dev/dsk/c1t500A09838350B481d32s0</pre>
6	Mount the file system that you will use when you copy the boot data. Example: The following example mounts the SAN file system. <pre>mount /dev/dsk/c1t500A09838350B481d32s0 /mnt/boot lun</pre>
7	Create the required directory structure on the bootable LUN and copy the boot data. Enter: <pre>ufsdump 0f - <source_boot_device> cd/ <mntpoint of bootable_ lun>; ufsrestore rf -)</pre> Note_____ If your configuration boots off more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun's document Sun StorEdge SAN Foundation Software 4.4 Configuration Guide for instructions. _____ Example: The following example copies the information from c0t0d0s0. <pre># ufsdump 0f - /dev/rdisk/c0t0d0s0 (cd /mnt/bootlun; ufsrestore rf -)</pre>

Step	Action																																																																																																		
8	Edit the vfstab file to on the SAN file system. Change the swap, root, and file systems in the vfstab file to show the boot device instead of the local device. One way to do this is to edit the file using vi: <pre>cd /mnt/bootlun/etc/ vi vfstab</pre> Example: The following shows the vfstab entry for the bootable LUN. <table><tr><td>#device</td><td>device</td><td>mount</td><td>FS</td><td>fsck</td><td>mount</td><td>mount</td></tr><tr><td>#to mount</td><td>to fsck</td><td>point</td><td>type</td><td>pass</td><td>at boot</td><td>options</td></tr><tr><td>#</td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>fd</td><td>-</td><td>/dev/fd fd</td><td>-</td><td>no</td><td>-</td><td></td></tr><tr><td>/proc</td><td>-</td><td>/proc proc</td><td>-</td><td>no</td><td>-</td><td></td></tr><tr><td>/dev/dsk/c1t500A09838350B481d32s1</td><td>-</td><td>-</td><td>-</td><td>swap</td><td>-</td><td>no</td></tr><tr><td>/dev/dsk/c1t500A09838350B481d32s0</td><td>-</td><td>-</td><td>/dev/rdsk/c1t500A09838350B481d32s0</td><td>-</td><td>-</td><td>-</td></tr><tr><td>/</td><td>ufs</td><td>1</td><td>no</td><td>-</td><td>-</td><td>-</td></tr><tr><td>/dev/dsk/c1t500A09838350B481d32s6</td><td>-</td><td>-</td><td>/dev/rdsk/c1t500A09838350B481d32s6</td><td>-</td><td>-</td><td>-</td></tr><tr><td>/globaldevices</td><td>ufs</td><td>2</td><td>yes</td><td>-</td><td>-</td><td>-</td></tr><tr><td>/devices</td><td>-</td><td>/devices</td><td>devfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>ctfs</td><td>-</td><td>/system/contract</td><td>ctfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>objfs</td><td>-</td><td>/system/object</td><td>objfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>swap</td><td>-</td><td>/tmp tmpfs</td><td>-</td><td>yes</td><td>-</td><td>-</td></tr></table>	#device	device	mount	FS	fsck	mount	mount	#to mount	to fsck	point	type	pass	at boot	options	#							fd	-	/dev/fd fd	-	no	-		/proc	-	/proc proc	-	no	-		/dev/dsk/c1t500A09838350B481d32s1	-	-	-	swap	-	no	/dev/dsk/c1t500A09838350B481d32s0	-	-	/dev/rdsk/c1t500A09838350B481d32s0	-	-	-	/	ufs	1	no	-	-	-	/dev/dsk/c1t500A09838350B481d32s6	-	-	/dev/rdsk/c1t500A09838350B481d32s6	-	-	-	/globaldevices	ufs	2	yes	-	-	-	/devices	-	/devices	devfs	-	no	-	ctfs	-	/system/contract	ctfs	-	no	-	objfs	-	/system/object	objfs	-	no	-	swap	-	/tmp tmpfs	-	yes	-	-
#device	device	mount	FS	fsck	mount	mount																																																																																													
#to mount	to fsck	point	type	pass	at boot	options																																																																																													
#																																																																																																			
fd	-	/dev/fd fd	-	no	-																																																																																														
/proc	-	/proc proc	-	no	-																																																																																														
/dev/dsk/c1t500A09838350B481d32s1	-	-	-	swap	-	no																																																																																													
/dev/dsk/c1t500A09838350B481d32s0	-	-	/dev/rdsk/c1t500A09838350B481d32s0	-	-	-																																																																																													
/	ufs	1	no	-	-	-																																																																																													
/dev/dsk/c1t500A09838350B481d32s6	-	-	/dev/rdsk/c1t500A09838350B481d32s6	-	-	-																																																																																													
/globaldevices	ufs	2	yes	-	-	-																																																																																													
/devices	-	/devices	devfs	-	no	-																																																																																													
ctfs	-	/system/contract	ctfs	-	no	-																																																																																													
objfs	-	/system/object	objfs	-	no	-																																																																																													
swap	-	/tmp tmpfs	-	yes	-	-																																																																																													
9	Umount the file system on the bootable LUN. <pre>umount /mnt/bootlun</pre>																																																																																																		

Step	Action
10	Run the <code>show-disks</code> command to get the correct device path name for your FC card. It is important to look for the device path associated with the LUN you used during the earlier <code>format</code> step. Due to the way the OBP sees hardware, the device may show up with either “disk” or “ssd” as the terminating characters and those characters may differ from what was seen in the format output. Example: As can be seen in the sample output below, the device uses "disk" to represent the disk device. <pre>ok show-disks a) /pci@1f,700000/scsi@2,1/disk b) /pci@1f,700000/scsi@2/disk c) /pci@1e,600000/ide@d/cdrom d) /pci@1e,600000/ide@d/disk e) /pci@1d,700000/emlx@2,1/fp@0,0/disk f) /pci@1d,700000/emlx@2/fp@0,0/disk q) NO SELECTION Enter Selection, q to quit:</pre> Using the device name you wrote down in Step 3, boot to the bootable LUN. Two things to keep in mind when you do this: ◆ You can specify the slice of the bootable LUN by using ":a" for slice zero, ":b" for slice 1, and so on. ◆ If you are using Emulex HBAs or StorageTek HBAs based on Emulex technology, replace the "ssd" part of the string with "disk". Examples: The following examples show command lines for QLogic HBAs and Emulex HBAs. Qlogic: <pre># boot /pci@1f,2000/QLGC,qlc@1/fp@0,0/ssd@w500a09838350b481,20:a</pre> Emulex: <pre># boot /pci@1d,700000/SUNW,emlxs@2/fp@0,0/disk@w500a09818350b481,20:a</pre>

Step	Action
11	Using the device name you wrote down in Step 3, set the boot-device environment variable and boot the system. Note _____ You can specify the slice of the bootable LUN by using ":a" for slice zero, ":b" for slice 1, and so on. _____ Example: This example sets the boot device and reboots the system. <pre>ok setenv boot-device /pci@1d,700000/emlx@2,1/fp@0,0/disk@w500a09838350b481,f4:a ok boot -r</pre>
12	Re-enable MPxIO to provide path redundancy. ◆ For Solaris 10 update 5 and later: <pre># stmsboot -D fp -e</pre> ◆ For Solaris 10 releases prior to update 5: <pre># stmsboot -e</pre>
13	Verify that you booted to the correct LUN by using the <code>df -k</code> command. In addition, configure the system crash dump for the new bootable LUN boot device using the <code>dumpadm -d</code> command: <pre>dumpadm -d /dev/rdsd/c1t500A09838350B481d32s1</pre>

File system dump, restore for an x86/x64 system with MPxIO

x86/x64 (AMD) with MPxIO systems:
Copying data from locally booted server

If you have an x86/x64 (AMD) system that is using MPxIO and you want to copy data from a locally booted server, complete the following steps.

Step	Action
1	Use the <code>mount</code> command to identify the current boot device. Example: The following mounts the current boot device. <pre># mount / on /dev/dsk/c1d0s0 read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=1980 000 on Tue Jan 2 12:50:22 2007</pre>
2	Configure the LUN for ◆ The root file system ◆ The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices. In addition, you might want to set up the swap slice, <code>/var</code>, and other file systems. Use the <code>format</code> command to choose the LUN. Write down the device path for the LUN so you can refer to it later. Example: In this example, you first run the <code>format</code> command. Based on its output, you determine that you want number 33, which has the device path <code>/scsi_vhci/disk@g60a9800056716436636f393164695969</code>. <pre>33. c4t60A9800056716436636F393164695969d0 <DEFAULT cyl 9787 alt 2 hd 255 sec 63> /scsi_vhci/disk@g60a9800056716436636f393164695969</pre>
3	Create the SAN file system that you will dump the current file system to. Example: The following command line creates a new file system. <pre>newfs /dev/dsk/c4t60A9800056716436636F393164695969d0s0</pre>

Step	Action
4	Mount the file system that you will use when you copy the boot data. Example: The following example mounts the file system. <pre># mount /dev/dsk/c4t60A9800056716436636F393164695969d0s0 /mnt/bootlun</pre>
5	Create the required directory structure on the bootable LUN and copy the boot data. Enter: <pre>ufsdump 0f - <source_boot_device> cd/ <mntpoint of bootable lun>; ufsrestore rf -)</pre> Note_____ If your configuration boots off of more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun's document <i>Sun StorEdge SAN Foundation Software 4.4 Configuration Guide</i> for instructions. _____ Example: The following example copies the information from c1d0s0. <pre># ufsdump 0f - /dev/rdsk/c1d0s0 (cd /mnt/bootlun; ufsrestore rf -)</pre>
6	Edit the vfstab file on the SAN file system. Change the swap, root, and file systems in the vfstab file to show the boot device instead of the local device. One way to do this is to edit the file using vi: <pre>cd /mnt/bootlun/etc/ vi vfstab</pre> Example: The following shows the vfstab entry for the bootable LUN. <pre>#device device mount FS fsck mount mount #to mount to fsck point type pass at boot options # fd - /dev/fd fd - no - - /proc - /proc proc - no - - /dev/dsk/c4t60A9800056716436636F393164695969d0s1 - - swap - no - /dev/dsk/c4t60A9800056716436636F393164695969d0s0/dev/rdsk/ c4t60A9800056716436636F393164695969d0s0 / ufs 1 no - /dev/dsk/c4t60A9800056716436636F393164695969d0s6/dev/rdsk/c4t60A98000567164366 36F393164695969d0s6 /globaldevices ufs 2 yes - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes -</pre>

Step	Action
7	Modify the /boot/solaris/bootenv.rc file to reflect the new bootpath parameter. Use the device name noted above to replace the bootpath. To identify the boot slices, use "...:a" at the end if your boot slice is slice 0, use "...:b" for slice1, and so on. Note_____ You can also see the device path by executing a list command: <pre>ls -al /dev/rdisk/c4t60A9800056716436636F393164695969d0s2</pre> Example: The following is an example of a modified file. <pre>... setprop bootpath /scsi_vhci/disk@g60a9800056716436636f393164695969:a ... # # Copyright 2005 Sun Microsystems, Inc. All rights reserved. # Use is subject to license terms. # #ident "@(#)bootenv.rc 1.32 05/09/01 SMI" # # bootenv.rc -- boot "environment variables" # setprop kbd-type US-English setprop ata-dma-enabled 1 setprop atapi-cd-dma-enabled 0 setprop ttyb-rtts-dtr-off false setprop ttyb-ignore-cd true setprop ttya-rtts-dtr-off false setprop ttya-ignore-cd true setprop ttyb-mode 9600,8,n,1,- setprop ttya-mode 9600,8,n,1,- setprop lba-access-ok 1 setprop prealloc-chunk-size 0x2000 setprop bootpath /scsi_vhci/disk@g60a9800056716436636f393164695969:a setprop console 'ttya'</pre>
8	Install GRUB on the new file systems. <pre>cd /boot/grub /sbin/installgrub stage1 stage2 /dev/rdisk/c3t500A09818350B481d32s0</pre> Note_____ The GRUB configuration may indicate the wrong partition information for the bootable LUN. It may indicate that the boot slice is slice 1. If that is the case, you may need to change it to slice0 in the menu.lst file.

Step	Action
9	Update the GRUB bootloader: bootadm update-archive -R /mnt/boot lun
10	Umount the file system on the bootable LUN. umount /mnt/boot lun
11	Shut down the system. sync;halt
12	Configure your HBA's bios to boot to the bootable LUN. Then configure your x86/x64 (AMD) system to boot to the HBA. See the documentation for your server and HBA for more details.

File system dump, restore for an x86/x64 system without MPxIO

x86/x64 (AMD)
without MPxIO
systems: Copying
data from locally
booted server

If you have an x86/x64 (AMD) system that is not using MPxIO and you want to copy data from a locally booted server, complete the following steps.

Step	Action
1	Use the <code>mount</code> command to identify the current boot device. Example: The following mounts the current boot device. <pre># mount / on /dev/dsk/c1d0s0 read/write/setuid/devices/intr/largefiles/logging/xattr/onerror=panic/dev=1980 000 on Tue Jan 2 12:50:22 2007</pre>
2	Configure the LUN for ◆ The root file system ◆ The file systems that will be copied to the bootable LUN. You may want to create slices and place the file systems on the slices. In addition, you might want to set up the swap slice, <code>/var</code>, and other file systems. Use the <code>format</code> command to choose the LUN. Write down the device path for the LUN so you can refer to it later. Example: This example uses a QLogic HBA. You first run the <code>format</code> command. Based on its output, you determine that you want number 33, which has the device path <code>/pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20</code>. <pre>33. c3t500A09818350B481d32 <DEFAULT cyl 9787 alt 2 hd 255 sec 63> /pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20</pre> Note _____ Emulex HBAs have "fc" in the device name. QLogic HBAs have "fp" in the device name. _____
3	Create the SAN file system that you will dump the current file system to. Example: The following command line creates a new file system. <pre>newfs /dev/dsk/c3t500A09818350B481d32s0</pre>

Step	Action																																																																																																		
4	Mount the file system that you will use when you copy the boot data. Example: The following example mounts the file system. <pre>mount /dev/dsk/c3t500A09818350B481d32s0 /mnt/bootlun</pre>																																																																																																		
5	Create the required directory structure on the bootable LUN and copy the boot data. Enter: <pre>ufsdump 0f - <source_boot_device> cd/ <mntpoint of bootable_ lun>; ufsrestore rf -)</pre> Note_____ If your configuration boots off of more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Sun’s document Sun StorEdge SAN Foundation Software 4.4 Configuration Guide for instructions. _____ Example: The following example copies the information from c0t0d0s0. <pre>ufsdump 0f - /dev/rdsk/c0t0d0s0 (cd /mnt/bootlun; ufsrestore rf -)</pre>																																																																																																		
6	Edit the vfstab file on the SAN file system. Change the swap, root, and file systems in the vfstab file to show the boot device instead of the local device. One way to do this is to edit the file using vi: <pre># cd /mnt/etc/ # vi vfstab</pre> Example: The following shows the vfstab entry for the bootable LUN. <table><tr><td>#device</td><td>device</td><td>mount</td><td>FS</td><td>fsck</td><td>mount</td><td>mount</td></tr><tr><td>#to mount</td><td>to fsck</td><td>point</td><td>type</td><td>pass</td><td>at boot</td><td>options</td></tr><tr><td>#</td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>fd</td><td>-</td><td>/dev/fd fd</td><td>-</td><td>no</td><td>-</td><td></td></tr><tr><td>/proc</td><td>-</td><td>/proc proc</td><td>-</td><td>no</td><td>-</td><td></td></tr><tr><td>/dev/dsk/c3t500A09818350B481d32s1</td><td>-</td><td>-</td><td>swap</td><td>-</td><td>no</td><td>-</td></tr><tr><td>/dev/dsk/c3t500A09818350B481d32s0</td><td>/dev/rdsk/c3t500A09818350B481d32s0</td><td></td><td></td><td></td><td></td><td>/</td></tr><tr><td>ufs</td><td>1</td><td>no</td><td>-</td><td></td><td></td><td></td></tr><tr><td>/dev/dsk/c3t500A09818350B481d32s6</td><td>/dev/rdsk/c3t500A09818350B481d32s6</td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>/globaldevices</td><td>ufs</td><td>2</td><td>yes</td><td>-</td><td></td><td></td></tr><tr><td>/devices</td><td>-</td><td>/devices</td><td>devfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>ctfs</td><td>-</td><td>/system/contract</td><td>ctfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>objfs</td><td>-</td><td>/system/object</td><td>objfs</td><td>-</td><td>no</td><td>-</td></tr><tr><td>swap</td><td>-</td><td>/tmp tmpfs</td><td>-</td><td>yes</td><td>-</td><td></td></tr></table>	#device	device	mount	FS	fsck	mount	mount	#to mount	to fsck	point	type	pass	at boot	options	#							fd	-	/dev/fd fd	-	no	-		/proc	-	/proc proc	-	no	-		/dev/dsk/c3t500A09818350B481d32s1	-	-	swap	-	no	-	/dev/dsk/c3t500A09818350B481d32s0	/dev/rdsk/c3t500A09818350B481d32s0					/	ufs	1	no	-				/dev/dsk/c3t500A09818350B481d32s6	/dev/rdsk/c3t500A09818350B481d32s6						/globaldevices	ufs	2	yes	-			/devices	-	/devices	devfs	-	no	-	ctfs	-	/system/contract	ctfs	-	no	-	objfs	-	/system/object	objfs	-	no	-	swap	-	/tmp tmpfs	-	yes	-	
#device	device	mount	FS	fsck	mount	mount																																																																																													
#to mount	to fsck	point	type	pass	at boot	options																																																																																													
#																																																																																																			
fd	-	/dev/fd fd	-	no	-																																																																																														
/proc	-	/proc proc	-	no	-																																																																																														
/dev/dsk/c3t500A09818350B481d32s1	-	-	swap	-	no	-																																																																																													
/dev/dsk/c3t500A09818350B481d32s0	/dev/rdsk/c3t500A09818350B481d32s0					/																																																																																													
ufs	1	no	-																																																																																																
/dev/dsk/c3t500A09818350B481d32s6	/dev/rdsk/c3t500A09818350B481d32s6																																																																																																		
/globaldevices	ufs	2	yes	-																																																																																															
/devices	-	/devices	devfs	-	no	-																																																																																													
ctfs	-	/system/contract	ctfs	-	no	-																																																																																													
objfs	-	/system/object	objfs	-	no	-																																																																																													
swap	-	/tmp tmpfs	-	yes	-																																																																																														

Step	Action
7	Modify the /boot/solaris/bootenv.rc file to reflect the new bootpath parameter. Use the device name noted above to replace the bootpath. To identify the boot slices, use "...:a" at the end if your boot slice is slice 0, use "...:b" for slice1, and so on. Note_____ You can also see the device path by executing a list command: <pre>ls -al /dev/rdisk/c3t500A09818350B481d32s2</pre> _____ Example: The following sets the bootpath parameter. <pre>... setprop bootpath /pci@0,0/pci10de,5d@e/pci1077,138@0,1/fp@0,0/disk@w500a09818350b481,20:a ...</pre>
8	Install GRUB on the new file systems. <pre>cd /boot/grub /sbin/installgrub stage1 stage2 /dev/rdisk/c3t500A09818350B481d32s2</pre> Note_____ The GRUB configuration may indicate the wrong partition information for the bootable LUN. It may indicate that the boot slice is slice 1. If that is the case, you need to change it to slice0 in the menu.lst file. _____
9	Update the GRUB bootloader: <pre>bootadm update-archive -R /mnt/bootlun</pre>
10	Umount the file system on the bootable LUN. <pre>umount /mnt/bootlun</pre>
11	Shut down the system. <pre>sync;halt</pre>
12	Configure your HBA's bios to boot to the bootable LUN. Then configure your x86/x64 (AMD) system to boot to the HBA.

About this information

The following copyright statements and licenses apply to open source software components that are distributed with various versions of the NetApp Host Utilities software products. Your product does not necessarily use all the open source software components referred to below.

Where required, source code is published at the following location:

<ftp://ftp.netapp.com/frm-ntap/opensource/>

Copyrights and licenses

The following components are subject to the Perl Artistic License:

- ◆ FTP.pm portions copyright © 1995-2004 Graham Barr <gbarr@pobox.com>. All rights reserved.
- ◆ Telnet.pm portions copyright © 1997, 2000, 2002 Jay Rogers. All rights reserved.

Perl Artistic License:

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

“Package” refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

“Standard Version” refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

“Copyright Holder” is whoever is named in the copyright or copyrights for the package.

“You” is you, if you’re thinking about copying or distributing this Package.

“Reasonable copying fee” is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

“Freely Available” means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes, and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
 - Place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as ftp.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 - Use the modified Package only within your corporation or organization.
 - Rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 - Distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 - Accompany the distribution with the machine-readable source of the Package with your modifications.
 - Accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package.

You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package.
7. C or perl subroutines supplied by you and linked into this Package shall not be considered part of this Package.
8. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
9. THIS PACKAGE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The HBA API component is subject to the SNIA Public License 1.1.

Portions copyright © 2002 International Business Machines Corporation. All rights reserved.

Portions authored by Benjamin F. Kuo, Troika Networks, Inc.

Portions copyright © 2001 Microsoft Corporation. All rights reserved.

SNIA Public License 1.1:

1. Definitions.

1.1 “Commercial Use” means distribution or otherwise making the Covered Code available to a third party.

1.2 “Contributor” means each entity that creates or contributes to the creation of Modifications.

1.3 “Contributor Version” means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.4 “Covered Code” means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.5 “Electronic Distribution Mechanism” means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.6 “Executable” means Covered Code in any form other than Source Code.

1.7 “Initial Developer” means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.8 “Larger Work” means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.9 “License” means this document.

1.10 “Licensable” means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.11 “Modifications” means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

- ❖ Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.
- ❖ Any new file that contains any part of the Original Code or previous Modifications.

1.12 “Original Code” means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.13 “Patent Claims” means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.14 “Source Code” means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor’s choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.15 “You” (or “Your”) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, “You” includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, “control” means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty

percent (50%) of the outstanding shares or beneficial ownership of such entity

2. Source Code License.

2.1 The Initial Developer Grant. The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third-party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense, and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using, or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2 Contributor Grant. Subject to third-party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense, and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2)

separate from the Contributor Version; 3) for infringements caused by: i) third-party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1 Application of License. The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2 Availability of Source Code. Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3 Description of Modifications. You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4 Intellectual Property Matters.

(a) Third Party Claims. If Contributor has actual knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail

that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter.

(b) Contributor APIs. If Contributor's Modifications include an application programming interface and Contributor has actual knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations. Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5 Required Notices. You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be most likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity, or liability terms You offer.

3.6 Distribution of Executable Versions. You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligation of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a

different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7 Larger Works. You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation. If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.
5. Application of this License. This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.
6. Versions of the License.

6.1 New Versions. The Storage Networking Industry Association (the “SNIA”) may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2 Effect of New Versions. Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by the SNIA. No one other than the SNIA has the right to modify the terms applicable to Covered Code created under this License.

6.3 Derivative Works. If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases “Storage Networking Industry Association,” “SNIA,” or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the SNIA Public License. (Filling in the name of the Initial Developer,

Original Code, or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY. COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR, OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1 This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within a reasonable time after becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2 If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as “Participant”) alleging that:

(a) such Participant’s Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

8.3 If You assert a patent infringement claim against Participant alleging that such Participant’s Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior

to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4 In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. **LIMITATION OF LIABILITY.** UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.
10. **U.S. GOVERNMENT END USERS.** The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.
11. **MISCELLANEOUS.** This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation

which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. **RESPONSIBILITY FOR CLAIMS.** As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.
13. **MULTIPLE-LICENSED CODE.** Initial Developer may designate portions of the Covered Code as “Multiple-Licensed”. “Multiple-Licensed” means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of this License or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.
14. **ACCEPTANCE.** This License is accepted by You if You retain, use, or distribute the Covered Code for any purpose.

EXHIBIT A—The SNIA Public License.

The contents of this file are subject to the SNIA Public License Version 1.1 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

This file itself is the license.

Software distributed under the License is distributed on an “AS IS” basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Contributor(s): _____.

The Net_SSLeay.pm component is subject to the Net SSLeay License.

NetSSLeay License:

http://search.cpan.org/src/SAMPO/Net_SSLeay.pm-1.25/README License and Copying: Copyright © 1996-2002 Sampo Kellomaki. All rights reserved.

Distribution and use of this module is under the same terms as the OpenSSL package itself (i.e. free, but mandatory attribution; NO WARRANTY). Please consult LICENSE file in the root of the OpenSSL distribution. While the source distribution of this perl module does not contain Eric’s or OpenSSL’s code, if you use this module you will use OpenSSL library. Please give Eric and OpenSSL team credit (as required by their licenses). And remember, you, and

nobody else but you, are responsible for auditing this module and OpenSSL library for security problems, backdoors, and general suitability for your application.

SSLey License copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape's SSL. This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)". The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgment: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)".

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR

CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. The licence and distribution terms for any publicly available version or derivative of this code cannot be changed; i.e., this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.

The Netkit component is subject to the GNU GPL 2.0 license.

Copyright © 1989, 1991 Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

GNU General Public License 2.0:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution, and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION, AND MODIFICATION

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute, or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients'

exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement, or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this

License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found:

One line to give the program’s name and a brief idea of what it does. Copyright (c)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w`. This is free software, and you are welcome to redistribute it under certain conditions; type `show c` for details.

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program ‘Gnomovision’ (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

The Open Hardware Certification component is subject to the Robin's Nest Software license.

Copyright © 1990-2008 Robin's Nest Software, Inc. All rights reserved.

Robin's Nest Software License:

This Software Provided By Robin's Nest Software Inc.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in the supporting documentation, and that the name of the author not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL HE BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE, OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

The Windows Installer XML (WiX) Toolset component is subject to the Common Public License.

Copyright © Microsoft Corporation. All rights reserved.

Common Public License:

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION, OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

b) in the case of each subsequent Contributor:

- i) changes to the Program, and
- ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution ‘originates’ from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor’s behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

“Contributor” means any person or entity that distributes the Program.

“Licensed Patents” mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

“Program” means the Contributions distributed in accordance with this Agreement.

“Recipient” means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import, and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the

rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third-party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

- a) it complies with the terms and conditions of this Agreement; and
- b) its license agreement:
 - i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;
 - ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental, and consequential damages, such as lost profits;
 - iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and
 - iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

- a) it must be made available under this Agreement; and
- b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor

who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor (“Commercial Contributor”) hereby agrees to defend and indemnify every other Contributor (“Indemnified Contributor”) against any losses, damages, and costs (collectively “Losses”) arising from claims, lawsuits, and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor’s responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL,

SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published,

Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

Index

Symbols

/etc/system file
 setting variables 58

A

A/A
 Active/Active paths 83
A/P-C
 Active/Passive Concurrent paths 83
active paths
 Active/Active (A/A) 83
Active/Passive Concurrent paths
 A/P-C 83
ALUA
 also known as TPGS 6
 Asymmetric Logical Unit Access 6
 checking on host 135
 checking on igroup 138
 required for storage system 6
Array Support Library
 see ASL 72
ASL
 components 73
 definition 72
 fresh install 74
 multipathing reports 83
Attach Kits
 previous kits 2

B

basic_config
 configuration script 3
 description 56
 example 69
 example with LPFC drivers 60
 options 68
 options (LPFC drivers) 59
 options (native drivers) 63
 parameters 67
boot data

 copying 179, 200, 204
bootblk
 installing 179

C

CD-ROM
 obtaining software 43
cfmode
 identifying 57
checklists
 installation tasks 11
commands
 fcinfo_hba -port 130
 luxadm -display 130
 luxadm -probe 130
 metastat 132
 stmsboot 131
 vxdctl enable 108
 vxdiskadm 108
configuring system parameters 56
create_binding.pl
 definition 150
 prerequisistes 152
 rules for using 150
 using 154

D

device names
 use show-disks 192
diagnostic utilities
 <switch>_info 140
 aix_info 140
 collecting Solaris information 143
 collecting storage system information 141
 collecting switch information 144
 executing 140
 filer_info 140
 switch_info 144
disk shelves
 powering up 139
documentation

- obtaining 47, 53
- drivers
 - downloading software (Emulex) 23
 - Sun drivers for Emulex 34

E

- EFI 96
 - creating LUNs 96
 - Solaris labeling type 96
- Emulex
 - changing HBA mode 162
 - changing HBA to Emulex mode 168
 - downloading HBA software 21, 23
 - driver software 20
 - HBA firmware versions 35, 37
 - Sun native drivers 34
- equipment
 - powering up 139
- examples
 - checking ALUA on host 135
 - checking ALUA on igroup 138
 - checking MPxIO 135

F

- FC Solaris Attach Kit
 - removing 49
- FC Solaris Host Utilities
 - completing installation 92
 - contents 2
 - documentation 47, 53
 - MPxIO stack 10
 - obtaining 43
 - SAN Toolkit 3, 43
 - software 43
 - Veritas DMP stack 9
- FCode
 - enabling compatibility (QLogic) 166
 - upgrading (Emulex) 36
 - upgrading (QLogic) 38
- filer_info
 - diagnostic script 3
- firmware
 - upgrading 26, 29, 32
 - upgrading (Emulex) 35

- version 35, 37
- format utility
 - using to label LUNs 106, 111

H

- HBA
 - checking 130
- host
 - checking ALUA 135
 - checking MPxIO 135
 - communicating with storage system 139
 - powering up 139
- host utilities
 - recommended parameters 67

I

- igroup
 - checking ALUA setting 138
 - definition 92
 - methods for creating 96

L

- LPFC drivers
 - basic_config options 59
 - changing HBA to SD mode 168
 - configuring parameters 57
 - discovering LUNs 103, 104
 - firmware versions 32
 - modifying sd.conf 102
- lpfc.conf file
 - setting variables 57
- lputil
 - creating persistent bindings 101
- LUNs
 - 2TB 96
 - best practices for creating 97
 - creating bootable LUN 172
 - discovering (LPFC drivers) 102, 103, 104
 - discovering (native drivers) 105
 - EFI 96
 - labeling 106, 111
 - methods for creating 96
 - optimal performance 96

- ostype 96
- partitioning bootable LUN 172
- solaris ostype 96
- solaris ostyles 96
- solaris_efi ostype 96
- solaris_efi ostyles 96
- tasks for creating on filer 96
- types 96

M

MPxIO

- checking on host 135
- enable on SPARC systems 131
- stack 10
- stmsboot 131
- Sun StorEdge SAN Foundation Software 1
- transparent paths 5
- using with sanlun 5

mpxio_set

- configuration script 3

multipathing

- displaying information 119
- MPxIO 5

N

native drivers

- configuring 62
- discovering LUNs 105
- Emulex 34
- getting controller number
 - HBAs
 - getting controller number 105
- setting variables 62
- upgrading FCode (Emulex) 36
- upgrading FCode (QLogic) 38
- upgrading firmware (Emulex) 35

NOW site

- downloading software 43

O

OpenBoot

- definition 183
- modifying 183

ostype

- solaris 96
- solaris_efi 96

P

parameters

- recommended 67

paths

- not visible 5

persistent binding

- creating 150
- creating with create_binding.pl 154
- creating with lputil 101
- definition 92, 100
- methods for creating 100

powering up equipment 139

Q

QLogic

- changing HBA mode 166
- enabling FCode compatibility 166
- upgrading FCode 38

R

removing

- attach kit 49
- FC Solaris Host Utilities 50

S

SAN boot

- changing Emulex HBA mode 162
- changing HBA to SD mode(LPFC) 168
- changing QLogic HBA mode 166
- configuration overview 161
- copying the boot data 179, 200, 204
- creating bootable LUN 172
- device names 192
- direct install 194
- dump, restore on SPARC 195
- dump, restore with MPxIO on x86/x64 (AMD)
 - processors 200
- dump, restore without MPxIO on x86/x64

- (AMD) processors 204
 - installing the bootblk 179
 - modifying OpenBoot 183
 - partitioning bootable LUN 172
 - performing direct install 178
 - prerequisites 160, 192
- SAN Toolkit 3
 - components 3
 - obtaining 43
 - san_version command 3
 - sanlun utility 3
- san_version
 - command 3
- sanlun
 - part of SAN Toolkit 3
 - paths not visible 5
 - syntax 122
 - using 115
- scripts
 - basic_config 3
 - filer_info 3
 - mpxio_set 3
- sd.conf file
 - configuring to discover LUNs 102
 - definition 93
- sd_set 68
- software
 - downloading 43
 - obtaining from CD-ROM 43
- Solaris
 - displaying HBA information 124
 - displaying LUN information 116
- Solaris Volume Manager
 - check version 132
 - SVM 132
 - works with MPxIO 6
- solaris_efi
 - LUN type 96
- SPARC processors
 - basic_config options 68
 - file system dump, restore 195
- ssd.conf file
 - setting variables 62
- ssd_set 68
- storage system

- communicating with host 139
 - powering up 139
- Sun StorEdge SAN Foundation Software 1
- switch_info
 - executing 144
- switches
 - powering up 139
 - verifying setup 138
 - zoning 139
- system parameters
 - configuring 56
 - configuring manually 57
 - pre-requisites 55, 66

T

- TPGS
 - also known as ALUA 6
 - Target Port Group Support 6
- troubleshooting
 - support information 15

U

- uninstall script 49

V

- Veritas
 - configuring LUNs 114
 - stack 9
 - volume manager 6
- Veritas Volume Manager
 - configuring LUNs 108
- VID/PID
 - setting 67
- volume manager
 - Solaris Volume Manager 6
 - Veritas Volume Manager 6

X

- x86/x64 (AMD) processors
 - basic_config options 68
 - file system dump, restore with MPxIO 200
 - file system dump, restore without MPxIO 204

Z
ZFS

volumes 96
zoning 139

