

FCP Solaris™ Host Utilities 4.0 for Veritas®

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 215-02097_A0
October, 2006

Copyright and trademark information

Copyright information

Copyright © 1994–2006 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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 Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “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 NETWORK APPLIANCE 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.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. 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 Network Appliance.

The product described in this manual may be protected by one or more U.S. 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, DataFabric, Data ONTAP, Decru, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, Decru DataFort, EdgeFiler, FlexClone, FlexShare, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NOW NetApp on the Web, ONTAPI, OpenKey, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor, StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

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

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Setup Overview	1
Chapter 2	Removing FCP Solaris Attach Kits	5
	Removing the previous FCP Solaris Attach Kit	6
Chapter 3	Installing Emulex Software	7
	Before you begin	8
	Installing HBA software	9
	Upgrading firmware	12
Chapter 4	Configuring the Solaris Host	15
	Installing the FCP Solaris Host Utilities	16
	Configuring system parameters.	18
	Setting parameters manually	19
	Using the basic_config command	21
	Removing the FCP Solaris Host Utilities.	24
Chapter 5	Completing the Configuration	27
	Configuration overview.	28
	Mapping igroups and LUNs	30
	Creating persistent bindings	31
	Configuring the host to discover LUNs	34
	Completing the configuration.	37
Appendix A	Using the create_binding.pl Script	39
	The create_binding.pl script	40

	Before you begin	42
	Using the create_binding.pl script	44
Appendix B	Configuring SAN Boot	49
	SAN boot overview.	50
	Configuration overview.	51
	Creating the bootable LUN.	52
	Copy bootblks and data.	57
	Modifying OpenBoot.	60
Appendix C	Changing cfmode to Single Image	67
	Understanding cfmode	68
	Completing pre-requisites	70
	Upgrading to single image mode	71
Appendix D	VERITAS and Array Support Library	77
	Overview of the Array Support Library	78
	Installing the ASL	79
	Removing the ASL	82
	Optimizing the VERITAS DMP Restore Daemon f	83
Appendix E	Using sanlun	85
	Displaying host LUN information	86
	Displaying information about path	88
	Displaying host HBA information	90
Appendix F	Using Diagnostic Utilities	93
	Using Diagnostic Utilities	94
	Collecting storage system information	95
	Collecting Solaris information	97

Collecting switch information 98

Index 101

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 *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*, available on the NetApp on the Web™ (NOW) site at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ 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 on a host. It assumes that you are familiar with how to configure and administer a storage system and use the storage system-licensed SAN protocol, FCP, to store and retrieve data on a storage system.

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (filers, FAS appliances, and NearStore systems) are all *storage systems*—also sometimes called filers or storage appliances.
- ◆ The term *type* means pressing one or more keys on the keyboard. The term *enter* means pressing one or more keys and then pressing the Enter key.

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</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.
<i>enter</i>	Used to mean pressing one or more keys and then pressing the Enter key.

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

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

FCP Solaris Host Utilities for Veritas overview

The FCP Solaris Host Utilities contain software tools and documentation that help you configure a Solaris system as a host component in a NetApp SAN environment. These include:

- ◆ Configuration tools that help you configure HBA and system files as well as help you create persistent bindings.
- ◆ The `san lun` utility that helps you manage LUNs and the host HBA
- ◆ Diagnostic scripts that provide diagnostic information about components in your configuration. Customer Support may ask for output from these scripts.
- ◆ Documentation that describes how to configure the Solaris host (this manual).

This *Installation and Setup Guide* describes how to set up the host and configure the Solaris HBA to access LUNs on the storage system. You can configure the Solaris host using the tools provided in the FCP Solaris Host Utilities, HBA provided tools, or a combination FCP Solaris Host Utilities and HBA provided tools.

This manual provides beginning to end configuration information for the FCP Solaris host, but does not include detailed information about how to setup your storage system. For complete information about configuring and using your storage system, see the *Data ONTAP® Block Access Management Guide* for your version of Data ONTAP. This guide is available from the Data ONTAP documentation section of the NOW web site at <http://now.netapp.com/NOW/knowledge/docs>.

Uninstalling previous versions

You must uninstall the previous version of the kit, if required.

See Chapter 2, “[Removing FCP Solaris Attach Kits](#),” on page 5 for additional information.

Setup procedure

The procedure for setting up the FCP Solaris host adapter and storage system follows this basic sequence.

Step	Action
1	Install the HBA driver and HBA management suite in accordance with HBA instructions. Upgrade the HBA firmware, if required. This involves: 1. Downloading the driver, driver management software, and firmware software from the vendor site 2. Uncompressing the packages and performing package add See Chapter 3, “Installing Emulex Software,” on page 7 for additional information.
2	Configure the Solaris host. This involves ◆ Installing the FCP Solaris Host Utilities (optional) ◆ Configuring storage specific variables in the /kernel/drv/lpfc.conf and /etc/system files. See Chapter 4, “Configuring the Solaris Host,” on page 15 for additional information.
3	Complete the installation. This involves 1. Mapping LUNs and igroups 2. Creating persistent bindings using lputil. 3. Configuring the host to discover LUNs 4. Labeling LUNs 5. Configuring Volume Management software See Chapter 5, “Completing the Configuration,” on page 27 for additional configuration information. Note _____ This chapter creates persistent bindings using lputil. To create persistent bindings using the NetApp provided create_binding.pl script, refer to Appendix A, “Using the create_binding.pl Script,” on page 39

Optional procedures

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

Task	Additional information
Configure SAN boot.	See Appendix B, “ Configuring SAN Boot ,” on page 49
Update the storage system cluster setting to single image mode.	See Appendix C, “ Changing cfmode to Single Image ,” on page 67.
Install the Array Support Library (ASL)	See Appendix D, “ VERITAS and Array Support Library ,” on page 77
Use sanlun.	See Appendix E, “ Using sanlun ,” on page 85
Run diagnostic scripts.	See Appendix F, “ Using Diagnostic Utilities ,” on page 93

Where to go for host and storage system information

The following table describes the documents on the NetApp on the Web (NOW) site at the time this *Installation and Setup Guide* was written. Use the links below to access information about host and storage system requirements, supported configurations, and cabling instructions.

If you need information about...	See the following NetApp product documentation...
The latest supported configurations	◆ <i>Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Configuring the storage system	◆ <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

If you need information about...	See the following NetApp product documentation...
Managing SAN storage on the storage system	◆ <i>Data ONTAP® Block Access Management Guide for FCP</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>FilerView® Online Help</i> - available from within the FilerView GUI
Verifying compatibility of storage system with environmental requirements	◆ <i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Troubleshooting, operational considerations, and post-release information	◆ FCP Solaris Host Utilities Release Notes at http://now.netapp.com/NOW/knowledge/docs/san/
Upgrading Data ONTAP	◆ <i>Data ONTAP® Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

Chapter overview

This chapter explains how to remove a version of the FCP Solaris Attach Kit, if one is installed on your system.

If your system does not include a previous version of the FCP Solaris Attach Kit, continue configuration by following the instructions in Chapter 3, “[Installing Emulex Software](#),” on page 7.

Topics in this chapter

This chapter includes the following topics:

- ◆ “[Removing the previous FCP Solaris Attach Kit](#)” on page 6

Removing the previous FCP Solaris Attach Kit

Version of the uninstall script to use

You must use the `uninstall` script from the FCP Solaris Attach Kit software to uninstall previous versions of the Attach Kit software.

Backing up system files

The uninstall script automatically creates a backup copy of the `/kernel/drv/lpfc.conf` and `sd.conf` files as part of the uninstall procedure. However, you may wish to create a backup copy of your own before you begin the uninstall.

Uninstalling the Attach Kit software

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

Step	Action
1	Ensure that you are logged in as root.
2	Locate the FCP Solaris Attach Kit 2.0 software. By default, the FCP Solaris Attach Kit is installed in <code>/opt/NTAPsanlun/bin</code> .
3	From the <code>/opt/NTAPsanlun/bin</code> directory, enter the following command to remove the existing software: <code>./uninstall</code>
4	At the Are you sure you want to uninstall lpfc and sanlun packages? enter y. Result: The <code>uninstall</code> script creates a backup copy of the <code>/kernel/drv/lpfc.conf</code> and <code>sd.conf</code> files to <code>/usr/tmp</code>. The backup files are named as follows: ◆ <code>lpfc.conf.save</code> ◆ <code>sd.conf.save</code> If a backup copy already exists, the install script prompts you to overwrite the backup copy.
5	Reboot the host with the reconfigure option (<code>reboot -- -r</code>).

Configuration overview

This chapter explains how to obtain and install the Emulex software. It also describes how to upgrade firmware.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Before you begin”](#) on page 8
- ◆ [“Installing HBA software”](#) on page 9
- ◆ [“Upgrading firmware”](#) on page 12

Before you begin

Complete pre-requisite tasks

This manual assumes that you have installed and configured your storage system, connected it to the switch or fabric, and installed and connected the HBA on the host.

The following describes the tasks you should have completed before you download and configure the Emulex software.

Configuration Task	Additional information
Configure the storage system ◆ Install Data ONTAP ◆ Install licenses ◆ Connect the storage system to the fabric or switch	◆ Index to Data ONTAP documentation at http://now.netapp.com/NOW/knowledge/docs/ontap ◆ The <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ Switch documentation at http://now.netapp.com/NOW/knowledge/docs/san/
Configure the Solaris host ◆ Verify OS version and patches for the FCP Solaris Host Utilities ◆ Install HBA hardware ◆ Verify cables to fabric or switch	◆ The <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml ◆ HBA documentation provided with the HBA you use ◆ Index to switch and cable documentation at http://now.netapp.com/NOW/knowledge/docs/san/

Installing HBA software

Obtaining the software

The Emulex software and documentation is available on the Network Appliance page of www.emulex.com/ts/index.html.

Emulex components

The Emulex software is a .tar file that includes the following software components:

- ◆ The Host Bus Adapter driver software
- ◆ The HBA management suit, which includes HBAnyware and lputil.

This document uses lputil 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.

Downloading Emulex components

The latest HBA firmware is available on the Network Appliance page of www.emulex.com/ts/index.html. Because your HBA may use a previous firmware version, and you may wish to download new firmware at the same time as you download the rest of the Emulex software.

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

Step	Action
1	On the Solaris host, create a download directory for the Emulex software. Enter: <code>mkdir /tmp/emulex</code>
2	To download Emulex driver and firmware, access www.emulex.com/ts/index.html and click the NetApp link. 1. Follow directions to download the driver software to the /tmp/emulex directory. 2. Follow directions to download Emulex firmware to the /tmp/emulex directory.

Step	Action
3	Change to the /tmp/emulex directory and extract the .tar files that contain the Emulex software. Enter: <pre>tar -xvf <solaris_version>-<driver_version>.tar</pre> Result: The .tar files are extracted to the default directory.
4	Extract the .tar files that contain Emulex applications and driver software. Example: The following commands extract Emulex driver and management software. <pre>tar -xvf <EmLxApps><version_number>-solaris.tar tar -xvf lpfc-<version_number>-sparc.tar</pre>
5	Unzip and extract the HBAware software. Example: The following commands unzip and extract HBAware software: <pre>gzip -d HBAware-<version_number>-sparc.tar.gz tar -xvf HBAware-<version_number>-sparc.tar .</pre>

Adding the Emulex packages

To add the Emulex packages, follow the steps below.

Step	Action
1	From the directory where you downloaded the Emulex software, enter: <pre>pkgadd -d .</pre>
2	At the Select packages you wish to process prompt, enter the numbers that correspond to the Emulex driver and driver management software. Result: The system displays the pkgadd installation script.

Step	Action
3	Respond as follows to the pkgadd prompts: 1. Rebuild manual pages database for section 7d [y,n,?] y 2. Use IP networking over Fibre Channel [y,n?] n (This option is not supported on FCP attached NetApp storage systems.) 3. Do you want to continue with the installation of <lpfc> [y,n,?] y 4. Do you want to continue with the installation of <HBAnyware> [y,n,?] y 5. Select package(s) you wish to process (or "all" to process all packages). (default:all) [?,??,q] q
4	Depending on your configuration, take one of the actions below. ◆ Update HBA firmware, if required —follow the instructions in “Firmware upgrade procedure” on page 13 ◆ Configure the host system — follow the procedures in Chapter 4, “Configuring the Solaris Host,” on page 15 Note_____ You must reboot the host with the reconfigure option (reboot -s -r) to save the work that you completed in this procedure. _____

Upgrading firmware

Determining the firmware version

Follow the steps below to determine if you need to upgrade Emulex firmware.

Step	Action
1	Activate the lputil tool. Enter: <code>/usr/sbin/lpfc/lputil</code> Result: The software displays the Main Menu.
2	At the Enter choice prompt that follows the Main Menu, enter 2 to display adapter information. Result: The software displays the Adapter Information Menu.
3	At the Enter choice prompt that follows the Adapter Information Menu, enter 2 to display adapter revision information. Result: The software displays a choice of adapters.
4	At the Enter choice prompt that follows adapter choices, select the adapter for which you want revision information. Result: The software displays the firmware version for the adapter you specify.
5	To exit lputil: 1. Enter 0 to continue. 2. At the Enter choice prompt after the Adapter Information, enter 0 to return to the Main Menu. 3. At the Enter choice prompt after the Main Menu, enter 0 to exit lputil.
6	Access the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config to determine the current firmware requirement.

Firmware upgrade procedure

Follow the steps below to upgrade your firmware version.

Note

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

Step	Action
1	Activate the lputil tool. Enter: <code>/usr/sbin/lpfc/lputil</code> Result: The software displays the Main Menu.
2	At the Enter choice prompt that follows the Main Menu, enter 3 to being the firmware upgrade. Result: The software displays the Firmware Maintenance Menu and prompts you for the hardware adapter.
3	At the Enter choice prompt, select the adapter for which you want to upgrade firmware. Result: The software displays the Firmware Maintenance Menu.
4	At the Enter choice prompt, select 1 to load the firmware image. Result: The software prompts you for the location of the Emulex firmware.
5	Enter the specification for the directory where you downloaded the Emulex firmware. Result: The software installs the firmware and displays the Firmware Maintenance Menu.
6	To exit lputil: 1. At the Enter choice prompt after the Firmware Maintenance Menu, enter 0 to return to the Main Menu. 2. At the Enter choice prompt after the Main Menu, enter 0 to exit lputil.

Chapter overview

This chapter explains how install the FCP Solaris Host Utilities software and configure the Solaris host.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Installing the FCP Solaris Host Utilities”](#) on page 16
- ◆ [“Configuring system parameters”](#) on page 19
- ◆ [“Setting parameters manually”](#) on page 20
- ◆ [“Using the basic_config command”](#) on page 22
- ◆ [“Removing the FCP Solaris Host Utilities”](#) on page 25

Installing the FCP Solaris Host Utilities

Before you begin

Configuring the Solaris host involves installing the FCP SolarisHost Utilities on the Solaris system and configuring host parameters.

The FCP Solaris Host Utilities provides configuration tools that help you:

- ◆ Configure host system parameters (`basic_config`)
- ◆ Create persistent bindings (`create_binding.pl`)
- ◆ Display host and storage system information (`sanlun`)

You can also use tools from the HBA Management Suite to perform these functions. If you wish to use HBA management suite tools, you do not need to install the FCP Solaris Host Utilities. To configure the host using HBA management tools, refer to “[Configuring system parameters](#)” on page 19.

Verify host system requirements

You must verify that your host system meets requirements before you begin the installation. Refer to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ for the latest hardware and software requirements.

Obtaining the Host Utilities software

The FCP Solaris Host Utilities software is available on the:

- ◆ NetApp NOW page
- ◆ CD-ROM

Host Utilities contents

The FCP Solaris Host Utilities software contains the following tools.

- ◆ `basic_config` - This is a command line tool that helps you configure the `/kernel/drv/lpfc.conf` and `/etc/system` file to NetApp required values.
- ◆ `sanlun` utility- This is a command utility that helps you manage LUNs and the host HBA
- ◆ Diagnostic scripts - These scripts provide diagnostic information about components in your configuration. Customer support may ask for output from these scripts.

- ◆ `create_binding.pl` - This is an interactive script that helps you configure persistent bindings for the Solaris host.

Installing the Host Utilities from NOW

To install the FCP Solaris Host Utilities software from NOW, complete the following:

Step	Action
1	Download or copy the <code>santoolkit_solaris_sparc_v1.0.0.tar.Z</code> file to a directory on your host. Enter: <pre>cp santoolkit_solaris_sparc_v1.0.0.tar.Z /<directory_name></pre>
2	Unzip the file. Enter: <pre>gunzip santoolkit_solaris_sparc_v1.0.0.tar.Z</pre> Result: The software unzips the tar.Z files
3	Untar the file. Change to the directory where you copied the file, and enter: <pre>tar -xvf santoolkit_solaris_sparc_v1.00.tar</pre> Result: The Host Utilities scripts are extracted to the default directory.
4	Add the packages that you extracted from tar file. Enter: <pre>pkgadd -d NTAPSANTool.pkg</pre> Result: The packages are added to the <code>opt/NTAP/SANToolkit/bin</code> directory.

Installing the Host Utilities from the CD-ROM

To install the FCP Solaris Host Utilities software from the CD-ROM, complete the following:

Step	Action
1	Login to the host system as root.
2	Mount the CD-ROM on the host system and change to the directory where the CD-ROM is mounted.

Step	Action
3	Install the Host Utilities package by entering the following command: <pre>pkgadd -d <attachkit_file>.pkg</pre>

Obtaining documentation

You can view and download documentation from <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Documentation also is available on the CD-ROM that came with your FCP Solaris Host Utilities 3.0.

Configuring system parameters

Configuration Methods

Configuring system parameters involves modifying the `/kernel/drv/lpfc.conf` and `/etc/system` files to recommended values. You can configure system parameters by:

- ◆ Manually editing the `/kernel/drv/lpfc.conf` and `/etc/system` files
- ◆ Using the `basic_config` command. This command is provided as part of the FCP Solaris Host Utilities 3.0.

Before you begin

Complete the following pre-requisites before you configure system parameters.

- ◆ Enter the `fcpx show cfmodes` command, on the storage system, to identify the storage system's `cfmode` setting. Required values in the host configuration files vary and depend on the `cfmode` setting.

If you plan to change the storage system `cfmode` to single image, first complete the procedure in Appendix C, “[Changing cfmodes to Single Image](#),” on page 67. Then follow the procedure below to configure system parameters.

- ◆ Create your own backup of the `/kernel/drv/lpfc.conf` and `/etc/system` file. The `basic_config` command automatically creates a backup for you, but does not allow you to revert more than one time.

Setting parameters manually

Procedure overview To set parameters manually you must:

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

These are described in the sections that follow.

Identifying the cfmode

You must enter the `fcpl show cfmode` command, on the storage system, to identify the storage system's cfmode setting. Required values in the host configuration files vary, and depend on the cfmode.

Setting lpfc.conf variables

To set /kernel/drv/lpfc.conf variables that match the storage system cfmode, identify the cfmode and enter one of the following:

- ◆ For Standby or Mixed (clustered storage systems)
 - ❖ `linkdown-tmo=30`
 - ❖ `nodev-tmo=120`
- ◆ For Single_image, Partner or Dual Fabric (clustered storage systems)
 - ❖ `linkdown-tmo=30`
 - ❖ `nodev-tmo=30`
- ◆ For Stand Alone (single storage system)
 - ❖ `linkdown-tmo=255`
 - ❖ `nodev-tmo=255`

To reset the topology parameter, enter one 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=32
- ◆ lun-queue-depth=16
- ◆ scan-down=0
- ◆ num-iocbs=1024
- ◆ num-bufs=1024
- ◆ dqfull-throttle-up-time=4
- ◆ dqfull-throttle-up-inc=2

Setting /etc/system variables

To set /etc/system variables that match the storage system cfmode, identify your cfmode and enter one of the following:

- ◆ For Standby or Mixed (clustered storage systems)
 - ❖ set sd:sd_io_time=0x78 (120)
- ◆ For Single_image, Partner, or Dual Fabric (clustered storage systems)
 - ❖ set sd_io_time=0x1e (30)
- ◆ For Stand-Alone (single storage systems)
 - ❖ set sd_io_time=0xff (255)

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

- ◆ set sd:sd_retry_count=5

Using the basic_config command

The basic_config command

The basic_config command is a configuration script that you can use to:

- ◆ Update items in the /kernel/drv/lpfc.conf to the values recommended for NetApp storage systems
- ◆ Identify your Solaris OS version, and set recommended values for that version in the /etc/system file
- ◆ Update timeout parameters in the /etc/system and /kernel/drv/lpfc.conf file, based on the cfmodes that you specify

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 files.

Configuring system parameters

To configure system parameters, follow the steps below.

Step	Action
1	To display existing values, 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

Step	Action
2	To configure /etc/system and /kernel/drv/lpfc.conf file items to the recommended setting, enter: basic_config -m <cfmode> where:-m describes the cfmode. Arguments are: ❖ single_image ❖ partner ❖ dual_fabric ❖ standby ❖ mixed ❖ stand_alone Result: The software displays recommended settings and prompts you for confirmation.
3	Enter y to confirm the changes and update the files. Example: In the example that follows, the first command displays the variables that the basic_config command will reset. The next command resets the variables in the /kernel/drv/lpfc.conf and /etc/system file to recommended values for a storage system that uses single image cfmode. basic_config -p <pre># Contents in the /etc/system file: # Tunable variables: set sd:sd_io_time=0x1e # Contents in the Emulex lpfc.conf file: # Tunable variables: lun-queue-depth=16; # Tunable variables: tgt-queue-depth=32; # Tunable variables: scan-down=0; # Tunable variables: linkdown-tmo=30; # Tunable variables: nodev-tmo=30; # Tunable variables: num-iocbs=1024; # Tunable variables: num-bufs=1024; # Tunable variables: topology=4; # Tunable variables: dqfull-throttle-up-time=4; # Tunable variables: dqfull-throttle-up-inc=2; </pre> basic_config -m single_image

Additional options The `basic_config` command has the following options.

Option	Description	Notes
-l	For direct attached configurations, sets the topology parameter in the <code>/kernel/drv/lpfc.conf</code> file to loop (topology=4).	The topology option in the <code>/kernel/drv/lpfc.conf</code> 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.
-s	For switched configurations, sets the topology parameter in the <code>/kernel/drv/lpfc.conf</code> file to switch (topology=2).	
-r	Reverts changes to the <code>/kernel/drv/lpfc.conf</code> and <code>/etc/system</code> files made by the <code>basic_config</code> script. This is done by copying the backup file to the current file.	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/lpfc.conf</code> and <code>/etc/system</code> file.
-v	Displays version information for the <code>basic_config</code> script.	

Removing the FCP Solaris Host Utilities

Removing the software

Use the following command to remove FCP Solaris Host Utilities software from the /opt/NTAP/SANToolkit/bin directory.

```
pkgrm NTAPSANTool
```


About this chapter This chapter describes what you need to do to complete the configuration.

Topics in this chapter This chapter includes the following topics:

- ◆ “[Configuration overview](#)” on page 28
- ◆ “[Mapping igroups and LUNs](#)” on page 30
- ◆ “[Creating persistent bindings](#)” on page 31
- ◆ “[Configuring the host to discover LUNs](#)” on page 34
- ◆ “[Completing the configuration](#)” on page 37

Configuration overview

Task overview

Completing the configuration involves the following tasks.

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. 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.	Refer to “ Creating WWPN bindings with lputil ” on page 32 for information about creating bindings with lputil. Refer to “ Using the create_binding.pl Script ” on page 39 for information about creating bindings with the create_binding.pl script.
3. 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 “ Configuring the host to discover LUNs ” on page 34 for additional information.
4. Discover LUNs	LUNs are discovered when you perform a reconfiguration reboot (reboot -- -r) on the host. You can also discover LUNs using the /usr/sbin/devfsadm command.	Refer to “ Discovering LUNs ” on page 35 for additional information.

Task	Discussion	Additional information
5. Label LUNs	Use the Solaris format utility to label the LUNs.	Refer to Solaris documentation for complete information about the format utility.
6. Configure Volume Management software	If your configuration includes VERITAS Volume Management software, you must use VERITAS tools and utilities to bring the new LUNs under VERITAS control.	Refer to VERITAS documentation for additional information.

Mapping igroups and LUNs

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.
- ◆ Using FilerView
This method provides a LUN Wizard that steps you through the process of creating and mapping new LUNs.
- ◆ 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.

Create the igroups and LUNS

For complete information about creating igroups and LUNs, refer to *Data ONTAP™ Block Access Management Guide for FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

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 FCP Solaris Host Utilities supports persistent bindings between a WWPN on the storage system and a target ID on the host. Bindings between a WWNN and a target ID are supported in standby mode.

Methods for creating persistent bindings

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

- ◆ Emulex Light Pulse Utility (lputil) — This is an installation and configuration script that loads automatically when you install the Emulex driver. Refer to “[Creating WWPN bindings with lputil](#)” on page 32 for information about creating persistent bindings using lputil.
- ◆ Emulex HBAnyware — The is an 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 FCP Solaris Host Utilities software.
 - ❖ To install the FCP Solaris Host Utilities software, refer to “[Installing the FCP Solaris Host Utilities](#)” on page 16
 - ❖ To configure persistent bindings using the `create_binding.pl` script, refer to “[Creating WWPN bindings with create_binding.pl](#)” on page 44.
- ◆ 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 lputil

Follow the steps below to create persistent bindings using `lputil`.

Caution

Previous versions of the FCP 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 `lputil` to create persistent bindings.

Step	Action
1	On the Solaris host, enter the following: <pre>/usr/sbin/lputil</pre> Result : The system activates the <code>lputil</code> program, and displays the Main Menu.
2	At the Enter choice prompt, select 5 to configure persistent bindings. Result : <code>lputil</code> shows the Persistent Bindings Menu and prompts you for a binding method.
3	At the Persistent Bindings Menu, select 5 to bind automapped targets. Result: <code>lputil</code> displays a list of adapters.
4	Select the number next to the adapter for which you will configure persistent bindings. Result: <code>lputil</code> displays the automapped WWPN and WWNN associated with the adapter you specify.
5	Press Enter or Return to continue. Enter yes at the following prompt if the WWPN names displayed in the previous step match the WWPN names for the initiator: <pre>Bind all auto-mapped targets? y</pre> Result: <code>lputil</code> displays Select Binding Method options. Note_____ The FCP Solaris Host Utilities supports WWNN bindings only on filers in standby mode.

Step	Action
6	To bind by port name, enter 2 . Result: lputil creates persistent WWPN bindings for the adapter you specify and displays the Persistent Bindings Menu.
7	To display the new bindings, enter 2 .
8	Repeat steps 3 through 7 to configure persistent bindings for additional adapters.
9	Quit lputil. 1. Enter 0 to return to the Main Menu. 2. Enter 0 to exit lputil.
10	Select one of the actions below: ◆ To continue configuration, complete the procedure in “ Modifying the sd.conf file ” on page 34 ◆ To quit lputil and save the changes you have made, reboot the host with the reconfigure option (reboot -- -r)

Configuring the host to discover LUNs

Modifying the sd.conf file

You must add an entry for each new target ID in the sd.conf file so the host system knows which disks to probe when the system is rebooted.

To modify the sd.conf file, follow the steps below.

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. You must create an entry for every LUN mapped to the host. Example: The following is an entry in the /kernel/drv/lpfc.conf file: fcplib-bind-wwpn="50a980010300e073:lpfc0t0, where: ◆ wwpn - 50a980010300e073 ◆ lpfc0 - the lpfc driver and instance ◆ t0 - the target number You must add a corresponding entry in the /kernel/drv/sd.conf file is as follows: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> 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. _____

Discovering LUNs

Follow these steps to discover new LUNs and validate that they are visible on the host.

Step	Action
1	To discover new LUNs, select one of the methods below: ◆ Reboot the host with the reconfigure option (<code>reboot -s -r</code>) ◆ Enter <code>/usr/sbin/devfsadm</code> Note— Occasionally the <code>/usr/sbin/devfsadm</code> 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.

Step	Action
2	To discover LUNs for the Solaris host: 1. Enter /usr/sbin/format. 2. At the format> prompt, enter p to display partition menu options 3. At the partition> prompt, enter p to print the partition table for Solaris host. Result: The system displays LUN information similar to the following. 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. c3t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1c,600000/lpfc@1/sd@0,0 5. c4t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1d,700000/lpfc@2/sd@0,0 Specify disk (enter its number):

Completing the configuration

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 activate the format utility, enter: 1. /usr/sbin/format 2. Follow screen directions to format or label LUNs. Result: The .format script displays available disks and configuration options.

Configuring Volume Management

If your configuration uses Volume Management software, you must configure the LUNs so they are under the control of the Volume Manager.

- ◆ Use the VERITAS **vxctl 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.

Appendix overview

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

You can also create bindings using `lputil` or `HBAnyware`.

- ◆ To create persistent bindings with `lputil`, refer to “[Creating WWPN bindings with lputil](#)” on page 32
- ◆ To create persistent bindings with `HBAnyware`, refer to `HBAnyware` documentation

Topics in this appendix

This appendix includes the following topics:

- ◆ “[The create_binding.pl script](#)” on page 40
- ◆ “[Before you begin](#)” on page 42
- ◆ “[Using the create_binding.pl script](#)” on page 44

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 automatically bind the WWPNs of a single storage system's active ports to a target ID and saves these bindings in the `/kernel/drv/lpfc.conf` file.

The FCP Solaris Host Utilities supports WWNNs only in standby mode. 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 WWNN 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 FCP 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 FCP cluster failover mode is set to mixed or dual_fabric, the script creates bindings for all available A and B ports.
- ◆ When the storage system's FCP 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.

When the script does not automate bindings

The `create_binding.pl` script does not automatically update WWPN bindings if you already have existing WWPN bindings saved in the `/kernel/drv/lpfc.conf` file. The script instead provides information about the bindings that it finds. You might have existing bindings in the following situations:

- ◆ You have one host connected to multiple storage systems or storage clusters, and you already used the script to create bindings for one storage system or one storage cluster.
- ◆ You replaced a failed HBA on the storage system and installed a new HBA.
- ◆ You have WWPN bindings from a previous configuration.

In these situations, you can use the `create_binding.pl` script to obtain information about available paths from the host to the storage system. The script detects which storage system ports are connected to host HBA ports and displays the WWPN bindings for these paths on the host console. You can add these WWPN bindings to the `/kernel/drv/lpfc.conf` file yourself. You can edit the file manually or use the `lputil` utility to add these bindings.

Note

You should always create a backup of the `/kernel/drv/lpfc.conf` file before you edit it manually or change it using a configuration script.

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 FCP Solaris Host Utilities is installed.

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

Refer to “[Installing the FCP Solaris Host Utilities](#)” on page 16 for information about downloading and installing the FCP 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 FCP 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 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. Use this option if you have multiple paths to the storage system, but have not yet 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</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 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. Example: The <code>create_binding.pl</code> script adds the following entry to the /kernel/drv/lpfc.conf file: <code>fcp-bind-wwpn="50a980010300e073:lpfc0t0</code>, where: ◆ <code>wwpn</code> - 50a980010300e073 ◆ <code>lpfc0</code> - the lpfc driver and instance ◆ <code>t0</code> - the target number The corresponding entry in the /kernel/drv/sd.conf file is as follows: <pre>name="sd" parent="lpfc" target=0 lun=0;</pre> 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 (reboot --r).

This appendix describes what you need to do to configure SAN boot on a Solaris host, and also provides a detailed example.

Topics in this appendix

Topics in this appendix include the following:

- ◆ [“SAN boot overview”](#) on page 50
- ◆ [“Configuration overview”](#) on page 51
- ◆ [“Creating the bootable LUN”](#) on page 52
- ◆ [“Copy bootblks and data”](#) on page 57
- ◆ [“Modifying OpenBoot”](#) on page 60

SAN boot overview

What is SAN boot

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

- ◆ Remove the hard drives from your servers and use the SAN for your boot needs, thereby eliminating the cost associated with maintaining and servicing hard drives
- ◆ Use the SAN for both storage and boot needs, thereby consolidating and centralizing storage

Contents of this appendix

This appendix provides a procedural example that uses a Solaris host to configure a SAN boot LUN and install a new OS image onto it. The procedure is intended as a detailed example for creating a bootable LUN that can be applied to most configurations. Other examples of bootable LUNs, such as configurations that boot multiple volumes or use diskless servers are not described in detail. Refer to Solaris documentation for details about additional configuration methods.

Configuration overview

Pre-requisites

The procedure and example in this chapter assume the following:

- ◆ The host and storage system are configured properly and use software and firmware that is supported by the FCP Solaris Host Utilities.
- ◆ The host operating system is installed on a local disk, and uses a UFS file system.
- ◆ The OpenBoot firmware (Bootcode) is downloaded and installed on the HBA.

The OpenBoot firmware is available on the Emulex site. Follow steps in “[Upgrading firmware](#)” on page 12 for details on how to load this firmware onto the HBA.

For information about host, storage system and firmware versions, refer to the *Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Configuration overview

Configuring a bootable LUN involves:

1. 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
2. Copying the bootblks and boot data
 - ❖ Install the bootblks onto the bootable LUN
 - ❖ Copy the boot data from the source disk onto the bootable LUN
3. Modifying OpenBoot
 - ❖ Verify the OpenBoot version on the adapter
 - ❖ Set the FCP topology to the bootable LUN
 - ❖ Bind the adapter and bootable LUN
 - ❖ Create an alias for the bootable LUN
4. Rebooting the system.

Creating the bootable LUN

Create the LUN

Use standard storage system commands and procedures to create and map a LUN assigned to Host ID 0. The example in this appendix uses host side LUN 0, (c3t0d0s0) as the bootable LUN.

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: 1. Activate the Solaris format utility. Enter /usr/sbin/format. Result: The system displays Available Disk Selections. 2. 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 filesystem in the /etc/vfstab file. Result: The system selects the Solaris boot disk, and displays partition options.
2	Display partition information for the current Solaris boot disk: 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 boot device.

Step	Action
3	Determine the size and layout of the partitions of the source boot volume. Make a note of the following: ◆ 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). Devices you must include are referenced in the vfstab file. 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. _____

Step	Action																																																						
	Example: The following shows the contents of the /etc/ vfstab file. Devices c1t0d0s1 and c1t0d0s0 (in bold) reference the source boot device, and must be replicated on the bootable LUN. <pre>root /mnt/etc 507-> vi 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/c1t0d0s1 - - swap - no - /dev/dsk/c1t0d0s0 /dev/rdsk/c1t0d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes - agate-be:/vol/local/sun4-sunos5 - /usr/local nfs - yes rw,bg,noquota,intr agate-be:/vol/local/share - /usr/local/share nfs - yes rw,bg,noquota,intr</pre> The next example shows the output of the print option for the current boot device on the Solaris host. Partitions that correspond to the vfstab file are displayed in bold. Total disk cylinders available: 14087 + 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>0 - 13081</td><td>63.48GB</td><td>(13082/0/0) - 133122432</td></tr><tr><td>1</td><td>swap</td><td>wu</td><td>13082 - 13503</td><td>2.05GB</td><td>(422/0/0) 4294272</td></tr><tr><td>2</td><td>backup</td><td>wm</td><td>0 - 14086</td><td>68.35GB</td><td>(14087/0/0) 143349312</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>	Part	Tag	Flag	Cylinders	Size	Blocks	0	root	wm	0 - 13081	63.48GB	(13082/0/0) - 133122432	1	swap	wu	13082 - 13503	2.05GB	(422/0/0) 4294272	2	backup	wm	0 - 14086	68.35GB	(14087/0/0) 143349312	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	0 - 13081	63.48GB	(13082/0/0) - 133122432																																																		
1	swap	wu	13082 - 13503	2.05GB	(422/0/0) 4294272																																																		
2	backup	wm	0 - 14086	68.35GB	(14087/0/0) 143349312																																																		
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	Select the bootable LUN: - Navigate to the Available Disk Selections screen: ❖ At the partition> prompt enter q ❖ At the format> prompt enter q Result: The system displays the Available Disk Selections screen. - At the specify disk prompt, enter the number that corresponds to the LUN mapped to host controller and Target 0. Note_____ Make a note of the identifying information for the LUN you select. You will use this information in “Modifying OpenBoot” on page 60, when you bind the bootable LUN to the adapter and create a boot alias. Example: The following example selects c3t0d0 as the bootable LUN. Output displayed in bold is saved by the system administrator for future reference. <pre>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. c3t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1c,600000/lpfc@1/sd@0,0 5. c4t0d0 <NETAPP-LUN-0.2 cyl 6142 alt 2 hd 16 sec 640> /pci@1d,700000/lpfc@2/sd@0,0 Specify disk (enter its number): 4 selecting c3t0d0 [disk formatted]</pre>
5	To display information about the bootable LUN: - At the format> prompt, enter p to display partition menu options. - 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 modify. Modify the partitions on the bootable LUN to match the host boot device.

Step	Action
7	Exit the format utility. 1. At the <code>partition></code> prompt, enter q. 2. At the <code>format></code> prompt, enter q.

Copy bootblks and data

- 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 shsun12 5.10 Generic_118822-18 sun4u sparc SUNW,Sun-Fire-V440</pre>
2	Write the bootblk. Enter: <pre>/usr/sbin/installboot /usr/platform/<directory_name>/lib/fs/ufs/bootblk /dev/rdisk/<device_name></pre> Example: The following example installs bootblks in the <code>/usr/platform/sun4u</code> directory on <code>/dev/rdsk/c3t0d0s0</code>. <pre>/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdsk/c3t0d0s0</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 filesystem.
 - ◆ 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: <pre>newfs <device></pre> Example: The following creates a USF filesystem at /dev/rdisk/c3t0d0s0. <pre>newfs /dev/rdisk/c3t0d0s0</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/c1t0d0s1 - - swap - no - /dev/dsk/c1t0d0s0 /dev/rdisk/c1t0d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes - agate-be:/vol/local/sun4-sunos5 - /usr/local nfs - yes rw,bg,noquota,intr agate-be:/vol/local/share - /usr/local/share nfs - yes rw,bg,noquota,intr</pre>
2	Mount the root directory of the bootable LUN. Enter: <pre>mount <filesystem_name> /<mountpoint></pre> Example: The following mounts the file system on the bootable LUN. <pre>mount /dev/dsk/c3t0d0s0 /mnt</pre>
3	Create the required directory structure onto the bootable LUN and copy the boot data. Enter: <pre>ufsdump 0f - <source_boot_device> cd/ <mntpoint of bootable_ lun>; ufsrestore rf -)</pre> Example: The following copies boot data from the source boot device to the bootable LUN. <pre>ufsdump 0f - /dev/dsk/c1t0d0s0 (cd /mnt; ufsrestore rf -)</pre>

Step	Action
4	Edit the vfstab file to reference the bootable LUN. Enter: cd /mnt/etc/ vi vfstab Example: The following shows the vfstab entry for the bootable LUN (bold). <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/c3t0d0s1 - - swap - no - /dev/dsk/c3t0d0s0 /dev/rdisk/c3t0d0s0 / ufs 1 no - /devices - /devices devfs - no - ctfs - /system/contract ctfs - no - objfs - /system/object objfs - no - swap - /tmp tmpfs - yes - agate-be:/vol/local/sun4-sunos5 - /usr/local nfs - yes rw,bg,noquota,intr agate-be:/vol/local/share - /usr/loca</pre>

Modifying OpenBoot

What is OpenBoot

OpenBoot is the firmware that the Emulex software uses to startup the system. OpenBoot firmware also includes the hardware level user interface that you use to configure the bootable LUN.

Modifying OpenBoot

Modifying OpenBoot involves:

- ◆ Validating the OpenBoot version
- ◆ Setting the topology (point-to-point or arbitrated loop) to the bootable LUN
- ◆ Binding the HBA to the bootable LUN
- ◆ Creating an alias for the bootable LUN. This alias substitutes for the device address (/pci@1c,600000/lpfc@1/sd@0,0 in the example that follows) during subsequent boot operations.

Complete the following to modify OpenBoot.

Step	Action
1	Activate the OpenBoot environment. Enter: 1. halt - to halt the Solaris Operating system Result: The OpenBoot program displays the OpenBoot prompt. 2. setenv auto-boot? false - to stop automatic reboot during OpenBoot configuration 3. reset-all - reset the host and activate your changes
2	Check the OpenBoot version that is currently installed on the host adapter. Enter: 1. select /<lpfc_device> - to select the host adapter 2. .version - to display the OpenBoot version For information about the correct firmware version, refer to the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/. 3. unselect-dev - to unselect the host adapter

Step	Action
3	Set the topology to the bootable LUN. - At the OpenBoot prompt, enter probe-scsi-all Result: The system displays all scsi devices - Select the adapter associated with the bootable LUN. Enter: ❖ select <boot_LUN_pathname> - Set the network topology. Enter one of the following: ❖ set-ptp - to set point to point ❖ set-fc-al - to set arbitrated loop. - Enter reset-all to reset the host and activate your changes.

Step	Action
	Example: The administrator in the following example displays scsi information, and selects lpfc@1. This is the adapter that accesses the bootable LUN 0 on WWPN 50a980030200f2cb. The topology is set to point-to-point. <pre> {1} ok probe-scsi-all /pci@1f,700000/scsi@2,1 /pci@1f,700000/scsi@2 Target 0 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 1 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 2 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 3 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB /pci@1d,700000/lpfc@2 Target none DID 11f00 WWPN 50a980010200f2cb LUN 0 Disk NETAPP LUN 0.2 /pci@1c,600000/lpfc@1 Target none DID 11f00 WWPN 50a980030200f2cb LUN 0 Disk NETAPP LUN 0.2 {1} ok select /pci@1c,600000/lpfc@1 {1} ok set-ntp Flash data structure updated. Signature 4e45504f Valid_flag 4a Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 4 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 {1} ok reset-all </pre>

Step	Action
4	Bind the host adapter to the bootable LUN. 1. Select the adapter associated with the bootable LUN. Enter: ❖ select <boot_LUN_pathname> 2. Enter probe-scsi-all to display the device ID (DID), LUN and target of the bootable LUN. 3. Bind the adapter and LUN. Values that you enter for the <device_id>,<LUN_id> and <target_ID >are in hex format. Enter: ❖ did <device_id> <LUN_id> <target_id> set-boot-id. 4. Enter reset-all to reset the host and activate your changes.

Step	Action
	Example: The following example binds the host adapter lpfc@1 to LUN 0. <pre> {1} ok probe-scsi-all /pci@1f,700000/scsi@2,1 /pci@1f,700000/scsi@2 Target 0 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 1 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 2 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB Target 3 Unit 0 Disk HITACHI HUS10733ASUN72G PA05 143374738 Blocks, 70007 MB /pci@1d,700000/lpfc@2 Target none DID 11f00 WPN 50a980010200f2cb LUN 0 Disk NETAPP LUN 0.2 /pci@1c,600000/lpfc@1 Target none DID 11f00 WPN 50a980030200f2cb LUN 0 Disk NETAPP LUN 0.2 {1} ok select /pci@1c,600000/lpfc@1 {1} ok did 11f00 0 0 set-boot-id Flash data structure updated. Signature 4e45504f Valid_flag ce Host_did 0 Enable_flag 5 SFS_Support 0 Topology_flag 4 Link_Speed_flag 0 Diag_Switch 0 Boot_id 0 Lnk_timer f Plogi-timer 0 LUN (1 byte) 0 DID 11f00 WWPN 0000.0000.0000.0000 LUN (8 bytes) 0000.0000.0000.0000 {1} ok reset-all </pre>

Step	Action
5	Create an alias for the bootable LUN. 1. Enter probe-scsi-all - to validate the configuration 2. Create the alias. Enter: nvalias <alias_name> /<boot_LUN_pathname> Example: The administrator in the following example creates an alias called “netappdisk”. nvalias netappdisk /pci@1c,6000000/lpfc@1/sd@0,0
6	Boot the LUN. Enter: boot <alias_name> Example: The following command boots the alias for /pci@1c,6000000/lpfc@1/sd@0,0. boot netappdisk

Appendix overview This appendix describes how to upgrade your host to single image mode.

Topics in this appendix

This appendix includes the following topics:

- ◆ [“Understanding cfmode”](#) on page 68
- ◆ [“Completing pre-requisites”](#) on page 70
- ◆ [“Upgrading to single image mode”](#) on page 71

Understanding cfmode

What is single image mode?

Single image is a new cfmode setting that is available on Data ONTAP 7.1. In single image mode, both systems in the cluster share the same World Wide Node Name (WWNN).

Single image mode provides the following advantages:

The host can access all LUNs through any target port on the cluster.

- ◆ Single image mode is supported on all NetApp storage systems and by all supported FCP hosts.
- ◆ There are no switch limitations. You can connect an NetAppstorage system in single image mode to any supported FCP switch.

Reasons for using single image mode

You might want to migrate your configuration to single image mode for the following reasons:

- ◆ To increase the number of paths to the cluster
You can increase the number of available paths in your configuration because all ports on each storage system are available to the host.
- ◆ To support heterogeneous configurations
For example, you upgraded an existing SAN configuration with Solaris hosts and you want to add HP-UX or AIX hosts to the configuration.

Host requirements

HP-UX and AIX hosts do not support failover when the storage systems in a cluster are in standby mode. If you want to add these hosts to an existing homogenous configuration that has Solaris-only hosts, you have to change the cfmode setting on the storage system to either single image or partner.

Solaris hosts in the same environment as HP-UX or AIX hosts must also have multipathing software installed when the cfmode setting is single image or partner.

Impact of changing the cfmode setting

Changing the cfmode setting of a cluster impacts the following components of your configuration:

- ◆ Access to LUNs
When you change the cfmode setting, you change the available paths between the host and the storage cluster. Some previously available paths

no longer available and some new paths become available. You must reconfigure every host that is connected to the cluster to discover the new paths. The LUNs might be accessible but must not be used until you reconfigure the host.

- ◆ Multipathing policy

If you have multipathing software in your configuration, changing the `cfmode` setting might also affect the multipathing policy because there are more paths between the host and the cluster.

- ◆ Zoning

Updating the filer to single image `cfmode` changes the WWNN of the cluster and the World Wide Port Names (WWPNs) of the target ports on each storage system. If the cluster connects to a switch by using soft zoning (zoning by WWPN), you must check the zoning setup and make the required updates to your switch.

- ◆ Switch configuration and cabling

In mixed and dual fabric modes, the target ports log into the fabric in loop mode. In single image mode, the target ports log into the fabric using point-to-point (PTP) mode. This means you might need to update your switch configuration when you migrate from these modes to single image mode so that the target connects to the fabric as point-to-point device. For additional information, see the *Data ONTAP® Block Access Management Guide for FCP* at

http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Completing pre-requisites

Before you begin You should complete the following pre-requisite tasks before you change your cfmode.

Task	What to do
1. Check system requirements	You must verify that your system uses the supported version of driver and volume management software. Refer to the latest configuration information, on the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml.
2. Save configuration information	You must record and save system configuration information before you begin. On the host, make a backup of all the volume groups that contain LUNs. These backups must contain up-to-date information for both the data within the volume group and the volume group configuration. Enter and save output from the following commands: ◆ <code>sanlun lun show</code> - display LUNs visible to the host. ◆ <code>sanlun lun show -p</code> - display LUNs and paths visible to the host If your LUNs are under VERITAS disk management, save the output from the following commands: ◆ <code>vxdisk list</code> - display disks under VERITAS control ◆ <code>vxdisk path</code> - display disks and paths under VERITAS control ◆ <code>vxprint -htg diskgroup_name</code> - display VERITAS volumes and status
3. Identify affected hosts and initiators	You must reconfigure and reboot any hosts and initiators that access LUNs on the storage systems you are upgrading. To identify affected hosts and initiators, enter the following command: ◆ <code>igroup show</code>
4. Schedule downtime	Hosts cannot access LUNs while you change cfmode. You must allow sufficient downtime to change the storage system cfmode, and reconfigure and reboot the host. Enter commands that quiesce host I/O and take host applications off-line.

Upgrading to single image mode

Procedure overview

To upgrade an FCP attached Solaris host and storage system to single image mode you must complete the following tasks:

1. Prepare the host — This involves stopping I/O and protecting the configuration by unmounting file systems and deporting VERITAS disk groups.
2. Upgrade the storage system to single image mode — This involves setting the cfmode to single image. You may also need to resolve LUN mapping or OS-type conflicts, if they occur.
3. Validate zoning — This involves checking zoning and resolving zone mapping issues.
4. Reconfigure the host — This involves generating persistent bindings and updating the sd.conf file. You must also reconfigure parameters in the /kernel/drv/lpfc.conf and /etc/system files to the values required by the storage system.
5. Reboot the host — This involves a reconfiguration reboot of the system. If your configuration uses VxVM volume managment, you must also enter additional commands to access VERITAS disk groups and file systems.

These are described in the sections that follow.

Preparing the host

To prepare the host for upgrade to single image mode, complete the following steps.

Step	Action
1	Open a window to the Solaris host
2	Enter commands that quiesce I/O and halt applications on each storage system that you want to upgrade.

Step	Action
3	If you are using VERITAS VxVM software, unmount the file systems and deport the associated VERITAS disk groups. Enter: 1. <code>/usr/sbin/umount /mountpoint</code> — to unmount the file system 2. <code>/usr/sbin/vxdg deport diskgroup_name</code> — to deport the VERITAS diskgroup
4	Verify that the VERITAS diskgroup is deported. Enter: <code>/usr/sbin/vxdisk -o alldgs list</code> Result: VERITAS displays exported disk groups in parentheses.

Upgrading the storage cluster to single image mode

To upgrade the storage cluster to single image mode, complete the following steps.

Step	Action
1	Open a window to the console of each storage system you want to upgrade.
2	Login at the console of the storage system(s) you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts. <code>lun config_check -S</code> Result: Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique to each node in the cluster. You use the <code>igroup</code> command to change igroup otypes, if required.
3	Switch to advanced privileges by entering the following command: <code>priv set advanced</code>
4	Stop the FCP service. On each storage system, enter: <code>fcp stop</code>

Step	Action
5	Set the cfmode to single_image. On each storage system, enter: <pre>fcpx set cfmode single_image</pre> If you see messages indicating mapping conflicts, then return to Step 2 to resolve the mapping conflicts.
6	Start the FCP service on each system in the cluster. On each system, enter the following command: <pre>fcpx start</pre>

Validating zoning and connectivity

Changes to WWNNs and WWPNNs: After you change the cfmode setting to single image both systems in the cluster share a single WWNN. One node in the cluster takes the WWNN of its partner. This means that the WWNN of one of the systems in the cluster has changed. The WWPNNs of the target ports also change. You might have to update your zoning configuration depending on the type of zoning you use. Soft zoning uses the WWPNN. Port zoning uses the physical location (switch port) in the zone. If you use soft zoning you must update your zones to use the new WWPNNs. Configurations that use hard zoning are not impacted.

Cabling and zoning: If you did not change your cabling, migrating to single image mode also changes the number of available paths to the LUNs. Some additional paths to the target become available because more target ports are visible to the host. You can use zoning or Data ONTAP portsets to reduce the number of paths in your configuration to comply with the number of paths supported by your multipathing software.

For details about validating and modifying switch zones see the documentation provided by your switch manufacturer.

For details about Data ONTAP portsets, see the [Data ONTAP 7.1 Block Access Management Guide](#).

You might also want to change your cabling if you have more paths between the host and the cluster and want to reduce the number of paths. It is recommended that after you change the cfmode setting, make sure each HBA has one primary and at least one secondary path to each LUN.

Reconfiguring the host

To reconfigure the host, complete the following steps.

Step	Action
1	Because changing the cfmode to single image changes the WWPN of the target HBA's, you must update the persistent bindings in the /kernel/drv/lpfc.conf file. You can use opt/NTAP/SANToolkit/bin/create_bindings.pl or lputil to re-generate persistent bindings in the /kernel/drv/lpfc.conf file. Select one of the options below. ◆ To use lputil enter: lputil From the lputil Main Menu, select 5 Follow screen prompts or refer to “Creating persistent bindings” on page 31 for additional information about using lputil. ◆ To use create_bindings.pl, enter: create_binding.pl -d -n storage_system_name Refer to Appendix A, “Using the create_binding.pl Script,” on page 39 for additional information about using create_binding.pl.
2	With a text editor, open the /kernel/drv/sd.conf file on the host and create an entry for any LUNs that are remapped. Note_____ You should create a backup of the sd.conf file before you make changes.
3	Update /etc/system and /kernel/drv/lpfc.conf parameters to the recommended values. Enter: basic_config -m single_image

Completing the upgrade

Follow the steps below complete the cfmode upgrade.

Step	Action
1	Perform a reconfigure reboot (reboot -- -r) on the host.
2	If your configuration uses a Volume Manager, you must import the VERITAS diskgroup and mount each file system. Enter: ◆ /usr/sbin/vxdg import <i>diskgroup_name</i> - to import the VERITAS diskgroup ◆ /usr/sbin/mount /mountpoint - to mount the file system

Appendix overview This appendix describes how to install and remove the Array Support Library (ASL). It also includes pointers to information about how to configure VERITAS for optimal performance during failover operations.

Topics in this appendix Topics in this appendix include the following:

- ◆ [“Overview of the Array Support Library”](#) on page 78
- ◆ [“Installing the ASL”](#) on page 79
- ◆ [“Removing the ASL”](#) on page 84
- ◆ [“Optimizing the VERITAS DMP Restore Daemon f”](#) on page 85

Overview of the Array Support Library

What the Array Support Library is

The Array Support Library (ASL) is an ONTAP qualified library that provides information about storage array attributes and multipathing configuration to the Device Discovery Layer (DDL) and Dynamic Multipathing (DMP) components of VERITAS Volume Manager™ (VxVM).

The DDL is a component of VxVM that discovers available enclosure information and the multipathing attributes of disks and disk arrays that are connected to a host system. The DDL calls ASL functions during the storage discovery process on the host.

You can add or remove ASLs from a running VxVM system. This way, dynamic multipathing support is added or removed from VxVM without stopping VxVM or rebooting the VxVM system.

What the ASL provides

The ASL provides the following information about SAN-attached filers:

- ◆ Enclosure-based naming—The name of the disk is based on the logical name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to DMP and VxVM about SAN-attached filers, instead of referring to them as JBOD devices or raw devices.

With enclosure-based naming, the name of the disk is based on the logical name of the enclosure, or disk array, in which the disk resides. The name of the disk is not based on the host assigned raw device name.

- ◆ Multipathing policy—Whether the storage is accessed as an Active/Active (A/A) disk array or an Active/Passive Group Concurrent (A/PG-C) disk array. The ASL also provides information about primary and secondary paths to 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 the ASL

ASL installation overview

Installing the ASL involves the following tasks:

- ◆ Determining the ASL version
- ◆ Verifying that your configuration meets system requirements.
- ◆ Installing a fresh version of the ASL *or* upgrading an existing ASL version.

Determining the ASL version

Use the VERITAS `vxddladm listversion` command to determine the version of the ASL that is installed on the host.

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

```
vxddladm listversion
LIB_NAME ASL_VERSION Min. VXVM version
=====
=
libvxp256.so vm-4.0-rev-1 4.0
libvxvpath.so vm-4.0-rev-1 4.0
libvxveritas.so vm-4.0-rev-1 4.0
libvxstorcomp.so vm-4.0-rev-1 4.0
libvxssa.so vm-4.0-rev-1 4.0
libvxshark.so vm-4.0-rev-1 4.0
libvxsenas.so vm-4.0-rev-1 4.0
libvxrdac.so vm-4.0-rev-1 4.0
libvxpurple.so vm-4.0-rev-3 4.0
libvxnec.so vm-4.0-rev-1 4.0
libvxlsiinf.so vm-4.0-rev-1 4.0
libvxhitachi.so vm-4.0-rev-1 4.0
libvxhds.so vm-4.0-rev-1 4.0
libvxfujitsu.so vm-4.0-rev-1 4.0
libvxemc.so vm-4.0-rev-2 4.0
libvxeccs.so vm-4.0-rev-1 4.0
libvxcscovrts.so vm-4.1-rev-1 4.0
libvxatf.so vm-4.0-rev-1 4.0
libvxap.so vm-4.0-rev-1 4.0
libnetapp.so vm-4.0-rev-1 4.0
#
```

Obtaining the ASL software

Check the VERITAS web site for the NetApp Solaris ASL version and follow directions.

If the NetApp Solaris ASL is not available, you can obtain it from the NetApp CD-ROM or the NetApp NOW Tools and Utilities site.

Installing ASL software from NOW or VERITAS sites

To install a fresh version of the ASL from either the VERITAS or NOW web sites, complete the following steps.

Step	Action
1	Log in to the VxVM system as the root user.
2	Check the VERITAS web site for the NetApp Solaris ASL version and follow directions. Note_____ If the NetApp Solaris ASL is not available, you can obtain it from the NetAppNOW Tools and Utilities. _____ Download the software and change to the directory where you downloaded the <code><asl_filename>.tar.Z</code> file.
3	Uncompress the file. Enter: uncompress <asl_filename>.tar.Z Example: The following command extracts an ASL file. <code>uncompress NTAPasl.4.0.2.tar.Z</code>
4	Extract the ASL software by entering the following command: tar -xvf <asl_filename>.tar Example: The following extracts ASL software. <code>tar -xvf NTAPasl.4.0.2.tar.Z</code>
5	Install the NTAPasl package by entering the following command: pkgadd -d <asl_filename>.pkg
6	Restart the vxconfigd daemon by entering the following command: vxctl enable

Step	Action
7	Verify the installation. ◆ If your host is connected to a storage system, you can verify the installation by entering the following command and then locating the Enclosure Type in the output of the <code>vxddmpadm listenclosure all</code> command. <pre> vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ===== OTHER_DISKS OTHER_DISKS OTHER_DISKS CONNECTED Disk Disk DISKS FAS960_0 FAS960 1032303 CONNECTED FAS9601 FAS960 1032305 CONNECTED </pre> ◆ If your host is not connected to a storage system, enter the following: <pre> vxddladm list support all LIBNAME VID PID ===== libvxap.so SUN All libvxatf.so VERITAS ATFNODES libvxeccs.so ECCS All libvxemc.so EMC SYMMETRIX libvxfujitsu.so FUJITSU GR710, GR720, GR730 libvxhds.so HITACHI All libvxhitachi.so HITACHI DF350, DF400, DF400F libvxlsiinf.so LSI INF-01-00 libvxnec.so NEC DS1200, DS1200F, DS3000SL libvxpurple.so SUN T300 libvxrdac.so VERITAS RDACNODES libvxsenas.so SENA All libvxshark.so IBM 2105 libvxssa.so SSA SSA libvxstorcomp.so StorComp OmniForce libvxveritas.so VERITAS All libvxxp256.so HP All libnetapp.so NETAPP All libvxvpath.so IBM VPATH_NODES </pre>

Installing ASL software from the CD-ROM

To install a fresh version of the ASL software from a CD-ROM, complete the following:

Step	Action
1	Log in to the VxVM system as the root user.

Step	Action
2	Mount the CD-ROM on the system and change to the directory where the CD-ROM is mounted.
3	Install the NTAPasl package by entering the following command: pkgadd -d <asl_filename>.pkg
4	Restart the vxconfigd daemon by entering the following command: vxctl enable
5	Reboot the host to activate your changes.

Upgrade Requirements

Only one version of the ASL that supports Data ONTAP is allowed on the host at any one time. However, you can run ASLs from other vendors on the same host.

Upgrading the ASL

To upgrade the ASL from a previous version, complete the following steps.

Step	Action
1	Verify that your system meets ASL requirements. Check the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml .
2	Log in to the host as root user.
3	Quiesce host I/O or take the host application offline.
4	Remove the previous version of the ASL by entering the following command: <pre>pkgrm <asl_filename></pre> Example: The following command removes a previous version of the ASL. <pre>pkgrm NTAPasl</pre> Note_____ You do not need to stop any volumes created on unaffected disk arrays, such as disk arrays from other vendors, before removing ASL support. This is also true for arrays or disks in the OTHER_DISKS category. OTHER_DISKS are local non-FCP-attached disks. Volumes created on these arrays remain accessible. _____
5	Follow direction in “ Obtaining the ASL software ” on page 80 to obtain and install new ASL software

Removing the ASL

Before you begin You do not need to stop any volumes created on unaffected disk arrays, such as disk arrays from other vendors, before removing ASL support. This is also true for arrays or disks in the OTHER_DISKS category. OTHER_DISKS are local non-FCP-attached disks. Volumes created on these arrays remain accessible because they are not multipathed.

Only one version of an ASL that supports Data ONTAP storage system can be installed on the host at any given time.

Removing the ASL To remove the ASL, complete the following steps.

Step	Action
1	Log into the host as the root user.
2	Enter the following command to remove the ASL package: pkgrm <asl_filename> Example: The following command removes a previous version of the ASL. pkgrm NTAPasl

Optimizing the VERITAS DMP Restore Daemon f

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 the frequency that the VERITAS daemon checks paths between the host and the storage system. Systems set to a longer polling interval take longer to detect path failures, and this can impact storage system performance during a failover operation.

Refer to the FCP Solaris Host Utilities Release Notes at <http://now.netapp.com/NOW/knowledge/docs/san/> for detailed information about recommended values.

This appendix describes how to use the `sanlun` utility to display information about the LUNs, paths to the LUNs, and the Solaris HBA.

Topics in this appendix

Topics in this appendix include the following:

- ◆ [“Displaying host LUN information”](#) on page 86
- ◆ [“Displaying information about path”](#) on page 88
- ◆ [“Displaying host HBA information”](#) on page 90

Displaying host LUN information

To display host LUN 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: <code>cd /opt/NTAP/SANToolkit/bin</code>

Step	Action
3	Enter the following command to display LUN information: <pre>./sanlun lun show [-v] [-d <i>host device filename</i> <i>all</i> <i>storage system name</i> <i>storage system name:storage system pathname</i>] -p</pre> -v produces verbose output. -d is the device option and can be one of the following: ◆ host device filename is the special device file name for the disk on Solaris (which might represent a storage system LUN). ◆ all lists all storage system LUNs under /dev/rdisk. ◆ storage system name lists all storage system LUNs under /dev/rdisk on a specified storage system. ◆ storage system name: storage system pathname lists all storage system LUNs under /dev/rdisk that are connected to the storage system path name LUN on a specified storage system. -p provides information about the primary and secondary paths available to the LUN when you are using multipathing. Result: The requested LUN information is displayed. Example: <pre># ./sanlun lun show -v -d /dev/rdisk/c2t1d2s2</pre> <pre>filer: lun-pathname device filename adapter lun size lun state filerX: /vol/vol1/hostA_lun2 /dev/rdisk/c2t1d2s2 lpfc1 150.0g (161061273088) GOOD</pre> <pre>Serial number: 2eFclndDjKlM Filer FCP nodename:50a9800055CC66DD Filer FCP portname:50a9800355CC66DD Filer adapter name: 5a Filer IP address: 10.22.110.36 Filer volume name:vol1 FSID:0xfffffc139 Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 LUN partition table permits multiprotocol access: no why: there can be only one data partition LUN has valid label: yes</pre> <pre>#</pre>

Displaying information about path

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: cd /opt/NTAP/SANToolkit/bin																																			
3	At the host command line, enter the following command to display LUN information: <pre>./sanlun lun show -p all</pre> -p provides information about the primary and secondary paths available to the LUN when you are using multipathing. all lists all storage system LUNs under /dev/rdsk. Example: <pre># ./sanlun lun show -p all filer1:/vol/solaris/lun0 (LUN 0) /dev/rdsk/c3t0d0s2 (c3t0d0s2) 300g (322122547200) lun state: GOOD</pre> <table><thead><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></thead><tbody><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c3t0d0s2</td><td>lpfc0</td><td>9a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>primary</td><td>/dev/rdsk/c4t0d0s2</td><td>lpfc1</td><td>11a</td><td></td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c3t2d0s2</td><td>lpfc0</td><td></td><td>9b</td><td>A/PG-C</td></tr><tr><td>enabled</td><td>secondary</td><td>/dev/rdsk/c4t2d0s2</td><td>lpfc1</td><td></td><td>11b</td><td>A/PG-C</td></tr></tbody></table>	path state	path type	device filename	host HBA	target port	partner port	Multipath policy	enabled	primary	/dev/rdsk/c3t0d0s2	lpfc0	9a		A/PG-C	enabled	primary	/dev/rdsk/c4t0d0s2	lpfc1	11a		A/PG-C	enabled	secondary	/dev/rdsk/c3t2d0s2	lpfc0		9b	A/PG-C	enabled	secondary	/dev/rdsk/c4t2d0s2	lpfc1		11b	A/PG-C
path state	path type	device filename	host HBA	target port	partner port	Multipath policy																														
enabled	primary	/dev/rdsk/c3t0d0s2	lpfc0	9a		A/PG-C																														
enabled	primary	/dev/rdsk/c4t0d0s2	lpfc1	11a		A/PG-C																														
enabled	secondary	/dev/rdsk/c3t2d0s2	lpfc0		9b	A/PG-C																														
enabled	secondary	/dev/rdsk/c4t2d0s2	lpfc1		11b	A/PG-C																														

Interpreting `sanlun lun show -p` output

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

- ◆ path state—Whether the path is enabled or disabled.
- ◆ path type—Whether the path to the LUN is primary, secondary, or standby.
 - ❖ 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 if the `cfmode` is either mixed or standby and 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 or FAS3000 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 FAS27 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:
 - ❖ A/A (Active/Active)—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.
 - ❖ A/PG-C (Active/Passive Group 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.
 - ❖ The primary and secondary controllers are each connected to a separate group of LUNS. If a single LUN in the primary controller's LUN group fails, all LUN's in that group fail over to the secondary controller's passive LUN group. This array type allows load balancing of I/O across multiple primary paths. Refer to VERITAS documentation for details about load balancing.

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: 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 FCP adapters. adapter name lists information for a specified adapter. Result: The FCP adapter information is displayed. Example: <pre>#./sanlun fcp show adapter -v lpfc0 WWPN:10000000c94ae06e WWNN: 20000000c94ae06e driver name: lpfc model: LP11002-M4 model description: Emulex LP11002 4Gb/s 2-Port FC : PCI-X Mode-2 HBA SFF Diag Optics serial number: BG53164700 hardware version: 1036406d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.10A7 (B2F2.10A7) Number of ports: 1 port type: Fabric port state: Operational supported speed: Unknown negotiated speed: 2 GBit/sec OS device name: /devices/pci@1d,700000/fibre-channel@1 lpfc1 WWPN:10000000c94ae06d WWNN: 20000000c94ae06d driver name: lpfc model: LP11002-M4 model description: Emulex LP11002 4Gb/s 2-Port FC : PCI-X Mode-2 HBA SFF Diag Optics serial number: BG53164700 hardware version: 1036406d driver version: 6.02h; HBAAPI(I) v2.0.e, 11-07-03 firmware version: 2.10A7 (B2F2.10A7) Number of ports: 1 port type: Fabric port state: Operational supported speed: Unknown negotiated speed: 2 GBit/sec OS device name: /devices/pci@1d,700000/fibre-channel@1,1</pre>

This appendix describes how to use the diagnostic utilities that are provided with the FCP Solaris Host Utilities.

Topics in this appendix

Topics in this appendix include the following:

- ◆ [“Using Diagnostic Utilities”](#) on page 94
- ◆ [“Collecting storage system information”](#) on page 95
- ◆ [“Collecting Solaris information”](#) on page 97
- ◆ [“Collecting switch information”](#) on page 98

Using Diagnostic Utilities

Diagnostic utilities

The FCP Solaris Host Utilities 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 95 for more information.
- ◆ `solaris_info`. See “[Collecting Solaris information](#)” on page 97 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 98 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 filer_info script tells you to delete or rename the file. Result: The output of various status commands are 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 storage_system_<<i>storage_system_name</i>>.tar.Z.
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: <pre>cd /opt/NTAP/SANToolkit/bin</pre>
2	At the command prompt, enter the following command: <pre>solaris_info [-d <directory>] [-n <file_name>] -f</pre> <code>-d <directory></code> specifies the location for the file that <code>solaris_info</code> creates. <code>-n <file_name></code> specifies the name for the output file and directory names. <code>-f</code> 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 <code>-f</code> 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 FCP Solaris Host Utilities software. A utility is provided for each switch that is supported for the 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 are 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 Emulexswitch is placed the brocade_switch_brcd1.tar.Z file.
3	Send the resulting file to Customer Support for analysis.

Index

Symbols

/etc/system file
variables 20

A

Array Support Library
see ASL 78

ASL

- components 78
- definition 78
- fresh install 79
- removing 82
- upgrade installation 81
- upgrade requirements 81

B

basic config command
description 21
options 23

boot data
copying 57

bootable LUN
creating 52
partitioning 52

bootblk
installing 57

C

configuring system parameters 21
create_binding.pl
definition 40
pre-requisites 42
rules for using 40
using 44

D

diagnostic utilities
<switch>_info 94

aix_info 94
collecting Solaris information 97
collecting storage system information 95
collecting switch information 98
executing 94
filer_info 94
switch_info 98
using 93

documentation

obtaining 17

driver software

- adding packages 10
- downloading 9
- obtaining 9

E

Emulex HBA

- adding software packages 10
- downloading software 9
- driver software 9
- firmware versions 12

F

FCP Solaris Host Utilities

- documentation 17
- installation overview 28
- installation procedure 16
- obtaining 16
- overview 1
- post installation tasks 2
- removing 6
- setup procedure 1
- software 16
- Solaris requirements 16

files

switch_info 98

firmware

- upgrading 13
- version 12

format utility

using to label LUNs 37

I

- igroup
 - definition 28
- installation
 - overview 1
- installing the FCP Solaris Host Utilities 16

L

- lpfc.conf file
 - variables 19
- lputil
 - creating persistent bindings 32
- LUNs
 - discovering 34
 - labeling 37

M

- multipathing
 - displaying information 88

O

- OpenBoot
 - definition 60
 - modifying 60

P

- persistent binding
 - creating 40
 - creating with create_binding.pl 44
 - creating with lputil 32
 - definition 28, 31
 - methods for creating 31
- post installation procedures 2

R

- removing previous attach kit versions 6
- requirements
 - Solaris host 16

S

- SAN boot
 - configuring 52
 - copying the boot data 57
 - definition 50
 - installing the bootblk 57
 - modifying OpenBoot 60
 - partitioning the bootable LUN 52
 - pre-requisites 51
- sanlun
 - syntax 89
 - using 85
- sd.conf file
 - configuring to discover LUNs 34
 - definition 28
- single image mode
 - changing 68
 - configuration overview 71
 - definition 68
 - preparing the host 71
 - pre-requisites 70
 - reconfiguring the host 74
 - Solaris requirements 68
 - upgrading the storage system 72
 - validating zoning 73
- Solaris
 - displaying HBA information 90
 - displaying LUN information 86
 - preparing for upgrade to single image mode 71
 - requirements 16
 - requirements for single image mode 68
- switch_info
 - executing 98
- system parameters
 - configuring 21
 - configuring manually 19
 - pre-requisites 18

U

- uninstall script 6

V

- VERITAS

configuring LUNs 37

Z

zoning

validating for single image mode 73

