

FCP HP-UX® Host Utilities 4.0.1 for Native OS and 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 215-02989_A0
January 2007

Copyright and trademark information

Copyright information

Copyright © 1994–2007 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, FAServer, FilerView, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, 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, EdgeFiler, FlexClone, FlexShare, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, 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.

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

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	vii
Chapter 1	Host Utilities Installation and Setup Overview	1
	FCP HP-UX Host Utilities Overview	2
	FCP HP-UX Host Utilities installation overview	5
	Where to go for more information	8
Chapter 2	Installing the Host Utilities	11
	Installing the FCP HP-UX Host Utilities 4.0.1	12
	Removing the FCP HP-UX Host Utilities SAN Toolkit or Attach Kit.	17
Chapter 3	Creating, Configuring, and Managing LUNs	21
	Overview of adding LUNs to HP-UX	22
	Preparing the storage system for HP-UX LUNs	23
	Creating igroups and LUNs.	24
	Configuring the LUNs	35
	Displaying LUN and HBA information	52
	Displaying host LUN information	53
	Displaying host HBA information	57
Chapter 4	Configuring SAN Booting.	59
	Overview of SAN booting procedures	60
	(LVM only) Creating the boot LUN using LVM.	62
	(LVM only) Migrating an OS disk using MirrorDisk/UX on PA-RISC systems	68
	(LVM only) Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems	75
	(LVM only) Testing the alternate boot path for PA-RISC systems	80
	(LVM only) Troubleshooting sanboot configurations with PA-RISC systems	81
	(Veritas only) Creating the boot LUN using vxcp_lvmroot	83

Chapter 5	Troubleshooting	87
	Troubleshooting checklists	88
	Correcting problems with Veritas	91
	Correcting volume group multipathing with ntap_config_paths	92
	Handling a takeover and giveback	94
	Recovering from a host reboot during a takeover.	95
	Enabling the host to identify more than 8 LUNs per target HBA	98
	Commands you can use to troubleshoot problems	104
	NetApp's diagnostic utilities	109
	General troubleshooting	114
 Appendix A	 Veritas and Array Support Library	 115
	Overview of the Array Support Library	116
	Installing the ASL	117
	Removing the ASL	119
	Obtaining multipathing information	120
	Understanding ASL error message types.	125
	 Index	 127

Preface

About this guide

This guide describes how to configure an HP-UX® 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 HP-UX versions supported.

Audience

This guide is for system installers and administrators who are familiar with the HP-UX operating system. It assumes that you understand how to configure and administer a filer and use the filer-licensed SAN protocol, FCP, to store and retrieve data on a filer.

This guide is for system installers and administrators who:

- ◆ Are familiar with the HP-UX operating system.
- ◆ Understand how to configure and administer a storage system.
- ◆ Understand how to use 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 and are also sometimes called filers or storage appliances.
- ◆ *igroup* refers to a collection of unique identifiers, either FCP WWPNs in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs.
- ◆ *LUN* (Logical Unit Number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Initiator* refers to the system component that originates an I/O command over an I/O bus or network.
- ◆ *Target* refers to the storage system that receives an FCP I/O request and serves the data requested by the initiator.

Command conventions

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

Formatting conventions

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

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fcctest <i>adaptername</i></code> command, you enter the characters “<code>fcctest</code>” 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.

What is in this guide...	What it means...
<i>Enter, enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Also used to mean pressing one or more keys and then pressing the Enter key.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.

Special messages

This guide contains the following types of special messages:

Note

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

Caution

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

How to send your comments

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

About this chapter

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

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“FCP HP-UX Host Utilities Overview”](#) on page 2
- ◆ [“FCP HP-UX Host Utilities installation overview”](#) on page 5
- ◆ [“Where to go for more information”](#) on page 8

FCP HP-UX Host Utilities Overview

Important information about Host Utilities and Attach Kits

The FCP HP-UX Host Utilities for Native OS and Veritas® provide the tools necessary to connect HP-UX hosts to NetApp storage systems. In the past, NetApp used the term FCP HP-UX Attach Kit to refer to kits containing these tools. By changing the name to Host Utilities, NetApp more clearly indicates that this kit provides utilities you can use when working with HP-UX hosts and NetApp storage systems.

The FCP HP-UX Host Utilities work the HP-UX Logical Volume Manager (LVM) and the Veritas Volume Manager (VxVM) to enable you to have different volume management and pathing solutions.

Note

The Host Utilities provide the tools to work with the host and its storage. They do not provide volume managers or pathing software.

NetApp provides all the software and documentation for the HP-UX Host Utilities under a single FCP HP-UX Host Utilities 4.0.1 heading. Component version numbers are not tied to the Host Utilities version number, so it is possible for each component to have a different version number. For example, at the time this document was produced, it supported the FCP Host Utilities 4.0.1 using version 3.1 of the SAN Toolkit. If you are concerned about the version numbers, you can confirm that you have the most current components by checking the Host Utilities version number in the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Overview of the FCP HP-UX Host Utilities

The FCP HP-UX Host Utilities 4.0.1 provide the tools and information necessary to enable a HP-UX host to connect to your storage systems.

The Host Utilities consist of:

- ◆ Utilities, scripts, and support files, which make up the SAN Toolkit.
- ◆ Documentation, including this setup guide, *Release Notes*, and a quick command reference guide.

For information on the supported stack, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. NetApp adds components on an ongoing basis. NetApp qualifies new HBAs, patches and other host components between Host Utilities releases. This online site is updated frequently to provide you with information about the current stack.

Software tools provided by the Host Utilities

The FCP HP-UX Host Utilities provide a SAN Toolkit. The SAN Toolkit contains scripts and commands you can use to display information about LUNs, add required settings, and to troubleshoot your system.

The SAN Toolkit contains the following components:

- ◆ `filer_info` script. This script collects information about the storage system.
- ◆ `hpux_info` script. This script collects configuration information about your host.
- ◆ `ntap_config_paths` script. When you use PVlinks, you can use `ntap_config_paths` to add missing paths to LVM volume groups and place them in the correct order.
- ◆ `san_version` command. This command displays the version of the SAN Toolkit that you are running.
- ◆ `sanlun` utility. You can use this utility to display information about LUNs on the storage system that are available to this host.
- ◆ Additional diagnostic scripts. The following scripts can also be used to collect troubleshooting information:
 - ❖ `brocade_info` collects information about Brocade switches installed in the network.
 - ❖ `cisco_info` collects information about Cisco switches installed in the network.
 - ❖ `mcdata_info` collects information about McData switches installed in the network.
 - ❖ `SHsupport.pm` is the support library for the `*info` scripts.
 - ❖ `Telnet.pm` support library for `*info` scripts.
- ◆ The man pages for `sanlun`, the diagnostic scripts, and the `ntap_config_paths` script.

Volume managers and multipathing

The FCP HP-UX Host Utilities for Native OS and Veritas manage the volumes using either HP-UX LVM or Veritas VxVM. In some cases you may have your host set up to use both LVM and VxVM.

LVM uses PVlinks to provide alternate paths if a problem causes the active path to go away. If you use PVlinks, you also need to use the `ntap_config_paths` script to ensure that all the paths are in the LVM volume group and that they are in the correct order.

VxVM uses Dynamic Multipathing (DMP) to provide multipathing. If you want to use VxVM to manage your LUNs, you must install the Symantec Array Support Library (ASL) for NetApp Storage.

Remove previous versions before installing Host Utilities

The FCP HP-UX Host Utilities 4.0.1 is a new product. If you currently have the FCP HP-UX Host Attach Kit 1.1 installed, you should remove it before you install the Host Utilities.

See “[Removing the FCP HP-UX Host Utilities SAN Toolkit or Attach Kit](#)” on page 17 for additional information.

FCP HP-UX Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

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

Pre-requisites

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

1. Set up your system:

- ❖ Host OS and appropriate updates
- ❖ HBAs
- ❖ Drivers
- ❖ Storage system and Data ONTAP
- ❖ Switches (if used)

See the *Compatibility and Configuration Guide for NetApp FCP and ISCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for current information about the system requirements.

2. Verify that your storage system is:

- ❖ Licensed correctly
- ❖ Running the FCP service
- ❖ Using the correct `cfmode` setting. See the *Compatibility and Configuration Guide for NetApp FCP and ISCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for information about supported cluster failover modes
- ❖ Configured to work with the target HBAs (if needed)
- ❖ Set up to work with the HP-UX host and the initiator HBAs
- ❖ Set up with working volumes and qtrees (if desired)

Note

It is recommended that you create volumes and LUNs with space reservations enabled.

3. If you're using a fabric connection, verify that the switch is:

- ❖ Set up correctly

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

Installation overview

Installing the Host Utilities involves the following tasks:

1. Get a copy of the Host Utilities' SAN Toolkit software package.
 - ❖ Get the compressed file containing the software package by either downloading it or copying it from the CD.
 - ❖ Extract the software package from the compressed file using the `gunzip` command.
2. Install the SAN Toolkit package, which is a standard HP-UX depot file:
 - ❖ Execute the `swinstall` command to install the file.
 - ❖ Log out and log back in to enable the updates that the installation script makes to the system search paths to take effect.
 - ❖ If you are using VxVM, install the Symantec Array Support Library (ASL) for NetApp storage systems to manage the LUNs.
3. Configure the LUNs by performing the following tasks:
 - ❖ Create at least one igroup and at least one LUN. The igroup contains the WWPNs of the host HBAs.
 - ❖ One way to create igroups and LUNs is to use the `lun setup` command.
 - ❖ Map the LUN to the igroup.
 - ❖ Use `ioscan` on the host to discover the LUN.
 - ❖ Create device special files (`/dev/dsk/` and `/dev/rdisk/`) for the newly discovered LUN.
 - ❖ Configure the LUN to be managed by LVM or VxVM.
 - ❖ Display information about the LUNs and HBA.
 - ❖ Use the `sanlun` command.

For detailed information on these tasks

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

- ◆ [“Installing the FCP HP-UX Host Utilities 4.0.1”](#) on page 12
- ◆ [“Creating igroups and LUNs”](#) on page 24
- ◆ [“Configuring the LUNs”](#) on page 35
- ◆ [“Displaying LUN and HBA information”](#) on page 52

In addition, the section [“Where to go for more information”](#) on page 8 provides pointers to additional documents and web sites that can help you with your system setup.

Optional procedures

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

Task	Additional information
Configure SAN boot. SAN boot uses a LUN as a boot device for an HP-UX host	See “Configuring SAN Booting” on page 59
Run diagnostic scripts. Normally, you would only run these scripts if you encountered a problem. These scripts collect information about your system and place it in compressed file that you can send to customer support.	See “NetApp’s diagnostic utilities” on page 109

Where to go for more information

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 guide was written. They provide information about host and storage system requirements, supported configurations, and cabling instructions.

If you need information about...	See the following NetApp product documentation...
Known issues, system requirements, and last minute updates	The <i>Release Notes</i> , which are available at http://now.netapp.com/NOW/knowledge/docs/san/
A summary of some of the commands you might use with the Host Utilities	The <i>Quick Reference</i> for the FCP HP-UX Host Utilities, which is available from the NOW site
The latest supported configurations	◆ The appropriate interoperability matrix <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
Supported Fibre Channel SAN topologies	<i>FCP Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
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>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>FilerView® Online Help</i> - available from within the FilerView GUI ◆ <i>Best Practices for Reliability: Ongoing Maintenance</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm
Verifying compatibility of storage system with environmental requirements	◆ <i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Upgrading Data ONTAP	◆ <i>Data ONTAP® Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
Installing and configuring the HBA in your host	◆ Refer to your HBA vendor documentation.
Your host operating system and using its features, such as LVM and VxVM	◆ Refer to your operating system documentation. You can download HP-UX manuals in PDF format from http://docs.hp.com/

About this chapter

This chapter describes procedures for installing the Host Utilities software package for an HP-UX host connected to a NetApp storage system running Data ONTAP with FCP. It also includes instructions for uninstalling the software.

Note

All installation procedures require root privileges. In addition, you must have the host and HBA set up before you install your software package.

Topics in this chapter

This chapter contains information on the following topics:

- ◆ [“Installing the FCP HP-UX Host Utilities 4.0.1”](#) on page 12
- ◆ [“Removing the FCP HP-UX Host Utilities SAN Toolkit or Attach Kit”](#) on page 17

Installing the FCP HP-UX Host Utilities 4.0.1

Before you begin

The FCP HP-UX Host Utilities 4.0.1 provides support for the HP-UX using LVM and using Veritas VxVM. The Host Utilities consists of the SAN Toolkit software package.

Before you install the SAN Toolkit software package:

- ◆ You must verify that your host system meets Host Utilities requirements. 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.
- ◆ If you currently have the FCP HP-UX Attach Kit or Host Utilities SAN Toolkit installed, you must remove it. See “[Removing the FCP HP-UX Host Utilities SAN Toolkit or Attach Kit](#)” on page 17.

Once your host is set up, you can install the SAN Toolkit software package. NetApp provides the SAN Toolkit as a standard HP-UX depot file. You can install it the same way you normally install an HP-UX depot file:

- ◆ Use the `gunzip` command to uncompress the files.
- ◆ Execute the `swinstall` command to install the NetApp software.

Note

You do not supply a path name to the installation directory to the `swinstall` command. The installation directory is determined by the depot file.

Actions performed during the installation

The installation performs the following actions:

- ◆ Checks your system to confirm that it meets the system requirements. The installation fails if your system does not meet the requirements.
- ◆ Installs NetApp's `sanlun` and `ntap_config_paths` commands and diagnostic tools in the appropriated directories. (You cannot specify which directories.)
- ◆ Adds paths to the NetApp SAN Toolkit utilities and man pages to the appropriate HP-UX search paths.

Obtaining the software

Copy the SAN Toolkit file from CD-ROM in your Host Utilities media kit or download it from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on your HP-UX host.

Note

To verify that you have the latest NetApp software package, check the *FCP/iSCSI Configuration Guide*, which is online at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Installing the SAN Toolkit software

To install the SAN Toolkit software, complete the following steps.

Step	Action
1	Enter the following command to uncompress the <code>santoolkit_hpux_3_1.depot.gz</code> file: <pre># gunzip santoolkit_hpux_3.1.depot.gz</pre> Result: The system places the extracted software in the directory to which you downloaded or copied the <code>santoolkit_hpux_3.1.depot.gz</code> file.

Step	Action
2	Use the <code>swinstall</code> command to install the software. You do not need to enter a pathname with this command. The depot file provides the installation directory information. Caution _____ The installation fails if your system does not meet the system requirements. _____ Result: This command runs an installation script that verifies the status of your HP-UX setup. If your system meets the requirements, this script installs the NetApp <code>sanlun</code> utility and diagnostic scripts in the <code>/opt/NetApp/santools/bin</code> directory. Example: The following output is similar to what you might see when you install the software for the first time. This example assumes that you uncompressed the files in the <code>/tmp/hpux</code> directory on an HP-UX host called <code>machine1</code>. <pre># swinstall -s /tmp/hpux/santoolkit_hpux_3.1.depot NetApp_santoolkit ===== 12/06/06 14:13:33 EST BEGIN swinstall SESSION (non-interactive) (jobid=machine1-0027) * Session started for user "root@machine1". * Beginning Selection * Target connection succeeded for "machine1:/". * Source: /tmp/hpux/santoolkit_hpux_3.1.depot * Targets: machine1:/ * Software selections: NetApp_santoolkit.command_itanium,r=3.1,fr=3.1,fa=HP-UX_B.11.22_IA NetApp_santoolkit.command_parisc,r=3.1,fr=3.1,fa=HP-UX_B.11.00_32/64 NetApp_santoolkit.man,r=3.1 NetApp_santoolkit.support_scripts,r=3.1 * Selection succeeded. * Beginning Analysis and Execution * Session selections have been saved in the file "/.sw/sessions/swinstall.last". * The analysis phase succeeded for "machine1:/". * The execution phase succeeded for "machine1:/". * Analysis and Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log machine1-0027 @ machine1:". ===== 12/06/06 14:13:50 EST END swinstall SESSION (non-interactive) (jobid=machine1-0027)</pre>

Updates to the system search paths

By default, the installation script updates the system search paths. It adds the following paths:

- ◆ /opt/NetApp/santools/bin and /opt/NetApp/santools/man to the system search path.
- ◆ /opt/NetApp/santools/bin to /etc/PATH.
- ◆ /opt/NetApp/santools/man to /etc/MANPATH.

You might need to log out and then log in again for these changes to take effect.

If you modified your startup to ignore /etc/PATH and /etc/MANPATH, you should manually add the santools paths to your search path.

Array Support Library

If you want to manage LUNs by using VxVM software, you must install the Symantec Array Support Library for NetApp storage systems.

The ASL is not distributed with the Host Utilities. You must download it from the Symantec Web site. For more information using the ASL, see “[Veritas and Array Support Library](#)” on page 115.

Removing the FCP HP-UX Host Utilities SAN Toolkit or Attach Kit

Removing the Host Utilities or an Attach Kit

Before you install a version of the Host Utilities, you should remove any previous versions of the software. If you currently have a version of the Host Utilities installed, you should remove the SAN Toolkit software. If you have a version of the HP-UX Attach Kit installed, you should remove that software. Once you have removed the software, you can install the Host Utilities.

When you remove this software, you use the `swremove` command. The following sections provide examples of removing the SAN Toolkit and the Attach Kit.

Uninstalling the SAN Toolkit or an Attach Kit

To uninstall the software, complete the following step.

Step	Action
1	Use the <code>swremove</code> command to remove the software. ◆ To remove the Host Utilities SAN Toolkit, enter the following command: <pre># swremove NetApp_santoolkit</pre> ◆ To remove an FCP HP-UX Attach Kit, enter the following command: <pre># swremove santools</pre> Example of removing the SAN Toolkit: The following example shows output you might see when you execute the <code>swremove</code> command to remove the FCP HP-UX Host Utilities SAN Toolkit from an HP-UX host called <code>machine1</code>. <pre># swremove NetApp_santoolkit ===== 12/06/06 14:16:10 EST BEGIN swremove SESSION (non-interactive) (jobid=machine1-0028) * Session started for user "root@machine1". * Beginning Selection * Target connection succeeded for "machine1:". * Software selections: NetApp_santoolkit.command_itanium,l=/opt/NetApp/santoolkit,r=3.1,fr=3.1,fa=HP-UX_B.11.22_IA</pre>

Step	Action
	<pre> NetApp_santoolkit.command_parisc,l=/opt/NetApp/santoolkit ,r=3.1,fr=3.1,fa=HP-UX_B.11.00_32/64 NetApp_santoolkit.man,l=/opt/NetApp/santoolkit,r=3.1 NetApp_santoolkit.support_scripts,l=/opt/NetApp/santoolki t,r=3.1 * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "machine1:". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "machine1:". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log machine1-0028 @ machine1:". ===== 12/06/06 14:16:13 EST END swremove SESSION (non- interactive) (jobid=machine1-0028) Example of removing an Attach Kit: The following example shows output you might see when you execute the swremove command to remove an Attach Kit from an HP-UX host called machine1. # swremove santools ===== 06/20/05 16:25:06 EST BEGIN swremove SESSION (non-interactive) (jobid=machine1-0125) * Session started for user "root@machine1". * Beginning Selection * Target connection succeeded for "machine1:". * Software selections: santools.command_itanium,l=/opt/NetApp/santools,r=1.0,fr= 1.1,fa=HP-UX_B.11.22_IA santools.command_parisc,l=/opt/NetApp/santools,r=1.0,fr=1 .1,fa=HP-UX_B.11.00_32/64 santools.man,l=/opt/NetApp/santools,r=1.0 </pre>

Step	Action
	<pre> santools.support_scripts,l=/opt/NetApp/santools,r=1.0 * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "machine1:". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "machine1:". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log machine1-0125 @ machine1:". ===== 06/20/05 16:25:21 EST END swremove SESSION (non- interactive) (jobid=machine1-0125) </pre>

About this chapter

This chapter provides the following information:

- ◆ Tasks required to add LUNs to systems with HP-UX hosts
- ◆ How to create initiator groups (igroups) and LUNs on the storage system
- ◆ How to configure LUNs and make them accessible by the host
- ◆ How to use the sanlun utility to manage host LUN and HBA information

Note

The procedures in this chapter require root privileges.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Overview of adding LUNs to HP-UX”](#) on page 22
- ◆ [“Preparing the storage system for HP-UX LUNs”](#) on page 23
- ◆ [“Creating igroups and LUNs”](#) on page 24
- ◆ [“Configuring the LUNs”](#) on page 35
- ◆ [“Displaying LUN and HBA information”](#) on page 52

Overview of adding LUNs to HP-UX

Tasks you perform on the storage system

You perform the following tasks on the storage system when you add a LUN to an HP-UX host:

1. If you have a cluster, make sure the storage system is correctly set up for adding HP-UX LUNs by verifying that the `cfmode` is correctly set.
2. Create an igroup using HP-UX as the operating system type for the host HBAs.
3. Create a LUN and map it to the igroup.

Tasks you perform on the host

You must perform the following tasks on the HP-UX host when you add a LUN:

1. Discover the newly created LUNs.
2. Create device special files (`/dev/dsk/` and `/dev/rdisk/`) for the newly discovered LUN.
3. Configure device special files to be managed by the HP-UX Logical Volume Manager (LVM) or Veritas Volume Manager (VxVM).

Note

To perform some of these tasks, you must use NetApp's `sanlun` utility. When you installed the SAN Toolkit on the host, the installation program placed the utility in `/opt/NetApp/santools/bin` and added it to the system search path (see [“Updates to the system search paths”](#) on page 15).

Preparing the storage system for HP-UX LUNs

Preparing the storage system

Before you create the igroups or LUNs, you must prepare the storage system by confirming that it has the correct FCP cluster failover mode (cfmode) setting. The cfmode value determines how a cluster handles a failover.

Check the cfmode setting on the storage system

To check the cfmode settings on your storage system, enter the `fcv show cfmode` command.

Example: The following command line checks the cfmode on the storage system.

```
filerA> fcp show cfmode  
fcv show cfmode: single_image
```

Caution

To determine which cfmodes are supported by the FCP HP-UX Host Utilities, see the *Compatibility and Configuration Guide for NetApp FCP and ISCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. If the cfmode is not set correctly, contact Customer Support. Do not change the cfmode.

Creating igroups and LUNs

Tasks to create and map LUNs on the storage system

Before you can complete your HP-UX setup with a NetApp storage system, you must perform the following tasks:

- ◆ Create an igroup on the storage system. An igroup is a collection of host WWPNs.
 - ❖ Each igroup should map to a host and should include all the initiators used by the host to access its data.
- ◆ Create one or more LUNs on the storage system and map the LUNs to an igroup.
 - ❖ Ensure that the LUNs are at the root of a qtree or a volume.
 - ❖ Do not put LUNs in the root volume. The default root volume is /vol/vol0.
 - ❖ Enable space reservations.
 - ❖ Disable scheduled auto snapshots.

For more information about creating and managing LUNs, see the *Data ONTAP Block Access Administration Guide*.

Creating an igroup

To create an igroup, complete the following steps.

Note

The examples used in these steps assume that the application is running on a single host, not a host cluster. They also assume that you have already attached the host and storage systems to fabric.

Step	Action
1	Log in as root on the host.

Step	Action
2	Get the igroup command you need to run on the storage system. To do this, enter the <code>sanlun fcp show adapter -c</code> command and then copy the command line it displays. You will need to enter this command line in step 4. sanlun fcp show adapter -c The <code>-c</code> option tells the <code>sanlun</code> utility to generate the output the storage system needs to create the igroup. Result: When you execute this <code>sanlun</code> command, it displays an <code>igroup create</code> command with the host's WWPNs appears on the screen. The igroup's name matches the name of the host. Example: Running <code>sanlun fcp show adapter -c</code> produces the following output on an HP-UX system called hostA. <pre>Enter this filer command to create an initiator group for this system: igroup create -f -t hpux "hostA" 50060b000010a700 50060b000010a606</pre> Because the name of the host is "hostA," this command line produces an igroup called "hostA."
3	On the host in a different session, use the <code>telnet</code> command to access the storage system.
4	Copy the <code>igroup create</code> command from Step 2. Example: If you use the information in Step 2, you would copy the following command line: <pre>igroup create -f -t hpux "hostA" 50060b000010a700 50060b000010a606</pre>
5	Paste the command on the storage system command line and press Enter. Result: The storage system creates the igroup and gives it the same name as the host name.

Step	Action
6	On the storage system command line, verify that the new igroup exists by entering the following command: <pre>igroup show</pre> Result: The storage system displays the newly created igroup with the host's WWPNs. Example: The following is output from the <code>igroup show</code> command displays. Note If you change the <code>ostype</code> attribute to <code>hpux</code> after you've created the igroup, you must manually set the <code>vsa</code> attribute, which HP-UX sets correctly by default. "Changing the vsa value" on page 98 contains information about how to do this. <pre>filerX> igroup show -v hostA hostA (FCP): OS Type: hpux Member: 50:06:0b:00:00:10:a7:00 Member: 50:06:0b:00:00:10:a6:06 filerX></pre> Example: The following example is for a direct-attached configuration. <pre>filerY> fcp show initiator Initiators connected on adapter 5a: Portname Group 50:06:0b:00:00:10:a7:00 hostA Initiators connected on adapter 5b: None connected. filerY></pre>

Creating a LUN using the interactive script

The `lun setup` command starts an interactive script that prompts you for information required to perform the following tasks:

- ◆ Create a LUN.
- ◆ Create a new igroup (or use the igroup you created in the previous section).
- ◆ Map the LUN to an igroup.

When a prompt displays a default in brackets [], you can press Enter to accept the default.

You can interrupt the `lun setup` command at any time by pressing Ctrl-C. Any information that you entered is discarded.

Note

You can also create a LUN using the `lun setup` command in command line mode. Instead of answering the prompts from the interactive script, you would supply those values as command line options using the following format:

```
lun create -s <size> -t <type> [-o noreserve] <lun_path>
```

To create LUNs on the storage system using the `lun setup` command interactive script, complete the following steps.

Step	Action
1	On the command line, start the LUN creation process by entering the following command: <pre>lun setup</pre> Result: The storage system starts an interactive script that prompts you for information needed to create a LUN. Example: This example shows the output that appears when you enter the <code>lun setup</code> command. <pre>filerX> lun setup This setup will take you through the steps needed to create LUNs and to make them accessible by initiators. You can type ^C (Control-C) at any time to abort the setup and no unconfirmed changes will be made to the system.</pre>
2	Enter y in response to the prompt: <pre>Do you want to create a LUN? [y]: y</pre>
3	Enter hpux to specify the host operating system: <pre>Multiprotocol type of LUN (image/solaris/windows/hpux/aix/linux) [image]: hpux</pre>

Step	Action
4	Enter the LUN path. Caution_____ Do not put LUNs in /vol/vol0 (the root volume). A LUN path must be absolute. A LUN can only reside in a volume or qtree root. For example, to create a LUN with the name "lun0" in the qtree root /vol/vol1/q0, specify the path as "/vol/vol1/q0/lun0" Enter LUN path: <i>lun_path</i> <i>lun_path</i> is the path name to the LUN. Result: The name of the LUN (the LUN path) is defined. Example: Some examples of a <i>lun_path</i> are /vol/vol1/mtree1/lun_sdata or /vol/vol1/hostA_0. Enter LUN path: /vol/vol1/hostA_0
5	Enter y or n to specify whether you want to use space reservations. Note_____ NetApp strongly recommends that you enable space reservations. A LUN can be created with or without space reservations being enabled. Space reservation guarantees that data writes to that LUN will never fail. Do you want the LUN to be space reserved? [y]: y Result: The LUN you are creating will use space reservations.

Step	Action
6	Enter the LUN size: Size for a LUN is specified in bytes. You can use single-character multiplier suffixes: b(sectors), k(KB), m(MB), g(GB) or t(TB). Enter LUN size: <code>x_x</code> <code>x_x</code> is the size you want the LUN to be. Result: The setup script accepts the value you enter here as the size for the LUN you are creating. The quantity of usable disk space by the host varies depending on the host's operating system and the application using the LUN. Example: An example of a size you might enter is 5g, meaning 5 GB of raw disk. Enter LUN size: 5g
7	Enter a comment string, if desired, or press Enter in response to the next prompt: You can add a comment string to describe the contents of the LUN. Please type a string (without quotes), or hit ENTER if you don't want to supply a comment. Enter comment string:_____ Note_____ If you press Enter and do not add a comment, you can add a comment later with the <code>lun comment</code> command. If you are using the FilerView® browser interface, a LUN description is the same as a comment. For more information about the <code>lun comment</code> command, see the Data ONTAP <code>na_lun (1)</code> man page. _____

Step	Action
8	Specify the igroup in one of the following ways: ◆ To use an existing igroup, enter that group's name. You can see the existing igroups by entering ? at the prompt. Go to step 12 to continue setting up a LUN. ◆ To create an igroup, enter a unique name for the group when prompted. Go to step 9 to continue setting up the LUN. ◆ To accept the default igroup (shown in brackets), press Enter. Go to step 12 to continue setting up a LUN. The LUN will be accessible to an initiator group. You can use an existing group name, or supply a new name to create a new initiator group. Enter '?' to see existing initiator group names. Name of initiator group [<i>default_igroup</i>]: <i>new_group_name</i> <i>new_group_name</i> is either the name of an existing igroup or the name of a group you want to create. A group name can be something similar to hpux_groupB or hostA. Result: The program either creates a new igroup or uses the existing group that you specify. Example 1: This example uses an existing igroup. This example is continued in step 12. <pre>Name of initiator group [hostA]: hostA An initiator group named "hostA" already exists. hostA (FCP) (ostype: hpux): 50:06:0b:00:00:10:a7:00 (logged in on: vtic) 50:06:0b:00:00:10:a6:06 (logged in on: vtic) Do you want to use this existing group? [n]: y</pre> Example 2: This example creates a new igroup called hpuxB. This example is continued in Step 9. <pre>Name of initiator group [hostA]: hpuxB</pre>

Step	Action
9	When you create an igroup (as illustrated in Example 2 in the previous step), you must specify setup information for the igroup. Press Enter to accept FCP as the protocol: Type of initiator group <i>new_group_name</i> (FCP/iSCSI) [FCP]: Result: Choosing FCP means that the initiator WWPNs in this igroup are using FCP to access the LUN. Note_____ If you created an igroup as shown in “Creating an igroup” on page 24, you would have already specified this information, so you would not see this prompt or any other prompts until step 12. _____

Step	Action
10	Enter comma-separated port names in response to the next prompt: A Fibre Channel Protocol (FCP) initiator group is a collection of initiator port names. Each port name (WWPN) is 16 hexadecimal digits, separated (only) by optional colon (:) characters. You can separate port names by commas. Enter '?' to display a list of connected initiators. Hit ENTER when you are done adding port names to this group. Enter comma separated portnames: Note Enter ? to see the list of initiator WWPNs connected to the target HBAs and associated igroups (if they exist). The sanlun fcp show adapter command runs on the host and lists the host's initiator HBAs and their WWPNs. Example: <pre>Enter comma separated portnames:? Initiators connected on adapter 5a: Portname Group 50:06:0b:00:00:10:a7:00 hostA 50:06:0b:00:00:10:a6:06 hostA Initiators connected on adapter 7a: None connected.</pre> Enter comma separated portnames:50:06:0b:00:00:10:a7:00 Enter comma separated portnames: In this example, the initiator HBA with WWPN 50:06:0b:00:00:10:a7:00 is added to the igroup you are creating (new_group_name).
11	Enter hpux in response to the next prompt: OS type of initiator group "new_group_name" (default/solaris/windows/hpux/aix/linux) : hpux Result: The operating system for the igroup is defined.

Step	Action
12	Enter the LUN ID for the LUN that is mapped to the igroup in response to the next prompt. The program supplies the next available LUN ID. (Enter ? to display the LUNs already in use.) Note Make sure you record the LUN ID associated with the <i>lun_path</i> and the igroup name (<i>new_group_name</i>). The LUN will be accessible to all the initiators in the initiator group. Enter '?' to display LUNs already in use by one or more initiators in group "<i>new_group_name</i>". LUN ID at which initiator group "<i>new_group_name</i>" sees "<i>lun_path</i>" [0]: Result: The LUN ID is assigned to the LUN (LUN path) and mapped to the igroup. Example: This example contains output you might see in response to this prompt. The LUN will be accessible to all the initiators in the initiator group. Enter '?' to display LUNs already in use by one or more initiators in group "hostA". LUN ID at which initiator group "hostA" sees "/vol/vol0/foobar" [0]: 3333 <pre> LUN Path : /vol/vol0/foobar OS Type : hpux Size : 2g (2147483648) Comment : test lun Initiator Group : hostA Initiator Group Type : FCP Initiator Group Members : 50:06:0b:00:00:10:a7:00 : 50:06:0b:00:00:10:a6:06 Mapped to LUN-ID : 3333 </pre>
13	Respond to the next prompt by entering y to accept the configuration or n to discard the configuration: Do you want to accept this configuration? [y]: Result: If you enter y, the LUN is mapped to the specified igroup and all changes to the system are committed. At this point, Ctrl-C cannot undo the changes.

Step	Action
14	Terminate this session or continue to create additional LUNs by responding to the final prompt: <pre>Do you want to create another LUN? [n]:</pre>

Verifying the mapped LUN

To verify the mapped LUN, complete the following step.

Step	Action
1	On the command line, display information about the mapped LUN by entering the following command: <pre>lun show -v <i>lun-path</i></pre> <i>lun-path</i> is the path to a specific LUN. Result: The information for the specified LUN is displayed. Example: The following example shows information about the LUN hostA. <pre>filerX> lun show -v /vol/vol1/hostA /vol/vol1/hostA 2g (2147483648) (r/w, online, mapped) Serial#: 0cABC1234ABC Share: none Multiprotocol Type: hpux Maps: hostA=3333 filerX></pre>

Configuring the LUNs

Getting the host to discover the new LUNs

- There are two ways to get the host to discover the new LUNs and manage them using LVM or VxVM:
- ◆ Reboot the host. The host automatically discovers the new LUNs when you reboot the host.
 - ◆ Complete the following steps. If rebooting the host is not a reasonable action for your system, perform the steps listed here. Doing this enables the host to discover the LUNs without requiring a host reboot.
 - ◆ To manage the LUN using LVM, see “[Managing the LUNs by using LVM](#)” on page 39.
 - ◆ To manage the LUN using VxVM, see “[Managing LUNs by using VxVM](#)” on page 45.

Note_____

This chapter provides information for using LVM and VxVM with NetApp LUNs. For more detailed instructions about working with LVM, see the HP-UX documentation. For more information about working with VxVM, see the Veritas documentation.

Step	Action
1	Log in as root on the host.

Step	Action
2	Perform an <code>ioscan</code> on the HP-UX host to see whether it discovers the newly created LUNs: <code>ioscan -fn -C disk</code> Example: In this example, the output from an <code>ioscan</code> command provides the following information: ◆ There are no device special files for the LUNs yet. ◆ The LUN is visible to the host and S/W state of each path to the LUN is CLAIMED, which means that the path is available. <pre># ioscan -fn -C disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 9 0/1/1/0/1/1.2.0 sdisk CLAIMED DEVICE HP 36.4GATLAS10K3_36_SCA /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0 disk 1 0/4/2/0/4/0.1.8.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 2 0/4/2/0/4/0.1.8.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 3 0/4/2/0/4/0.1.12.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 4 0/4/2/0/4/0.1.12.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 5 0/7/1/0/4/0.1.8.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 6 0/7/1/0/4/0.1.8.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 7 0/7/1/0/4/0.1.12.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 8 0/7/1/0/4/0.1.12.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN</pre>

Step	Action
3	Create the device entries on the HP-UX host by entering the following command: <pre>ioinit -i</pre> If the <code>ioinit -i</code> command doesn't create device special files, use the <code>insf -e</code> command to install the device special files. Note_____ Do not use these disk device special files yet. They are not yet under the control of a volume manager. _____ Example: The following is an abbreviated example of the type of output you might see when you run the <code>ioinit -i</code> command. <pre># ioinit -i insf: Installing special files for sdisk instance 1 address 0/4/2/0/4/0.1.8.0.0.0.2 insf: Installing special files for sdisk instance 2 address 0/4/2/0/4/0.1.8.0.0.0.4 insf: Installing special files for sctl instance 6 address 0/4/2/0/4/0.1.8.255.0.0.0 insf: Installing special files for sctl instance 5 address 0/4/2/0/4/0.1.11.255.0.0.0 insf: Installing special files for sdisk instance 3 address 0/4/2/0/4/0.1.12.0.0.0.2 insf: Installing special files for sdisk instance 4 address 0/4/2/0/4/0.1.12.0.0.0.4 insf: Installing special files for sctl instance 7 address 0/4/2/0/4/0.1.12.255.0.0.0 insf: Installing special files for sctl instance 4 address 0/4/2/0/4/0.1.15.255.0.0.0 insf: Installing special files for sdisk instance 5 address 0/7/1/0/4/0.1.8.0.0.0.2</pre>

Step	Action
4	Display information about the device special files by entering the following command: <pre>sanlun lun show -p all</pre> For each LUN, write down which paths are primary; you need this information later in these steps. Example: The following example shows the type of output you see using the <code>sanlun lun show</code> command before you configure the LUN by using a volume manager. <pre># sanlun lun show -p all filer-16:/vol/vol1/lun_04 (LUN 4) 36g (38654705664) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: None Multipath-provider: None ----- host filer primary partner path path /dev/dsk host filer filer state type filename HBA port port ----- up secondary /dev/dsk/c13t0d4 fcd1 0d up primary /dev/dsk/c10t0d4 fcd1 0c up secondary /dev/dsk/c7t0d4 fcd0 0d up primary /dev/dsk/c4t0d4 fcd0 0c</pre>

Managing the LUNs by using LVM To configure the LUNs as disks by using the Logical Volume Manager (LVM), complete the following steps.

Step	Action
1	Create a physical volume on the LUN by entering the following command: <pre>pvcreate /dev/rdisk/path_disk_device</pre> <i>path_disk_device</i> is the path to the disk device that contains the LUN. Note_____ You only need to perform this process once for each LUN, even if you have multiple paths to that LUN. _____ Example: The <code>pvcreate</code> command produces output similar to the following: <pre># pvcreate /dev/rdisk/c10t0d4 Physical volume "/dev/rdisk/c10t0d4" has been successfully created.</pre>
2	Check the minor numbers of all existing volume groups by entering the following command: <pre>ls -l /dev/*/group</pre> The LVM requires a unique minor number for each volume group device entry. Example: The <code>ls -l /dev/*/group</code> command provides output similar to the following. In this example, the next unused minor device number is 0x010000. <pre># ls -la /dev/*/group crw-r----- 1 root sys 64 0x000000 Dec 11 19:24 /dev/vg00/group</pre>
3	Use the <code>mkdir</code> command to create a <code>/dev</code> entry: <pre>mkdir /dev/directory_name</pre> Example: This example continues the previous example and specifies a directory name that ends with the number '01', which is the next minor device number. Using the same minor number as the volume group number is an HP-UX convention; it is not a requirement. <pre># mkdir /dev/vg_ntap01</pre>

Step	Action
4	Create a device node on the host for the physical volume group using the <code>mknod</code> command. You must supply a unique minor device number. Don't enter a number that another volume group is using. Note _____ The device node you create is used only for volume group and logical volume operations; it is not used for I/O. _____ Example: This example creates a device node using the unique minor number 0x01000. <pre>mknod /dev/vg_ntap01/group c 64 0x010000</pre>
5	Use the <code>vgcreate</code> command to create a volume group with one primary path. Use a regular or block disk device node. Do not use raw disk devices. vgcreate <i>vg_name path_to_disk</i> <i>vg_name</i> is a volume group name, with or without the <code>/dev/</code> prefix. Example: In this example, the <code>vgcreate</code> command sets up a primary path to the volume group. <pre># vgcreate /dev/vg_ntap01 /dev/dsk/c10t0d4 Increased the number of physical extents per physical volume to 9215. Volume group "/dev/vg_ntap01" has been successfully created. Volume Group configuration for /dev/vg_ntap01 has been saved in /etc/lvmconf/vg_ntap01.conf</pre>

Step	Action
6	Add the remaining paths to the volume group using the <code>ntap_config_paths</code> utility, which is in the <code>/opt/NetApp/santools/bin/</code> directory. Enter the following command: <pre>ntap_config_paths</pre> Note_____ This utility adds missing NetApp paths to the LVM volume group and places them in the correct order. It specifies the disk device special files that correspond to the primary paths first and then the device special files that correspond to the secondary paths. For more information about the <code>ntap_config_paths</code> utility, see “Correcting volume group multipathing with ntap_config_paths” on page 92. Example: The <code>ntap_config_paths</code> utility adds the rest of the paths using the <code>vgextend</code> command. You can verify the path ordering with the <code>vgdisplay -v volume_group</code> command. <pre># ntap_config_paths Getting information from sanlun... Adding missing path with: vgextend /dev/vg_ntap01 /dev/dsk/c4t0d4 Volume group "/dev/vg_ntap01" has been successfully extended. Volume Group configuration for /dev/vg_ntap01 has been saved in /etc/lvmconf/vg_ntap01.conf Adding missing path with: vgextend /dev/vg_ntap01 /dev/dsk/c13t0d4 Current path "/dev/dsk/c4t0d4" is an alternate link, skip. Volume group "/dev/vg_ntap01" has been successfully extended. Volume Group configuration for /dev/vg_ntap01 has been saved in /etc/lvmconf/vg_ntap01.conf Adding missing path with: vgextend /dev/vg_ntap01 /dev/dsk/c7t0d4 Current path "/dev/dsk/c4t0d4" is an alternate link, skip. Current path "/dev/dsk/c13t0d4" is an alternate link, skip. Volume group "/dev/vg_ntap01" has been successfully extended. Volume Group configuration for /dev/vg_ntap01 has been saved in /etc/lvmconf/vg_ntap01.conf #</pre>

Step	Action
7	Create one or more logical volumes in the volume group using the following command line: <pre>lvcreate [-L -l] <size> <volume_group></pre> <size> with an uppercase -L option is the total size of the logical volume in megabytes. <size> with a lowercase -l option is the number of physical extents that you want to assign to the logical volume. The value you supply for <size> must not equal or exceed the available free physical extents in the volume group. The LUN reserves an area for physical volume, volume group, and LVM information. The <code>lvcreate</code> and the <code>lvextend</code> commands fail if the size you specify here does not fit into the volume group. You can check the available free physical extents by using the <code>vgdisplay -v volume_group</code> command. Note _____ After you create a logical volume, PVlinks path management is in effect on the logical volume. Example: The following example on the next page creates a logical volume for <code>/dev/vg_ntap01</code> that uses all the physical extents available for that volume. The <code>vgdisplay</code> command displays the number of available physical extents (511). The <code>lvcreate</code> command with a lowercase -l option assigns them to the logical volume. See the <code>man</code> pages for more information about <code>vgdisplay</code> and <code>lvcreate</code>.

Step	Action
	<pre> # vgdisplay /dev/vg_ntap01 --- Volume groups --- VG Name /dev/vg_ntap01 VG Write Access read/write VG Status available Max LV 255 Cur LV 0 Open LV 0 Max PV 16 Cur PV 1 Act PV 1 Max PE per PV 1016 VGDA 2 PE Size (Mbytes) 4 Total PE 511 Alloc PE 0 Free PE 511 Total PVG 0 Total Spare PVs 0 Total Spare PVs in use 0 # # lvcreate -l 511 /dev/vg_ntap01 Logical volume "/dev/vg_ntap01/lvol1" has been successfully created with character device "/dev/vg_ntap01/rlvol1". Logical volume "/dev/vg_ntap01/lvol1" has been successfully extended. Volume Group configuration for /dev/vg_ntap01 has been saved in /etc/lvmconf/vg_ntap01.conf # </pre>

Step	Action
8	(Optional) If you want to create a file system on the logical volume device special file, use the <code>newfs</code> command. <code>newfs -F vxfs /dev/character_device</code> <i>character_device</i> is the character device listed in the output of the <code>vgdisplay -v volume_group</code> command. Example: The following example creates a file system on the logical volume <code>/dev/vg_ntap01</code>: <pre># newfs -F vxfs /dev/vg_ntap01/rlvol1 version 4 layout 2093056 sectors, 2093056 blocks of size 1024, log size 1024 blocks unlimited inodes, largefiles not supported 2093056 data blocks, 2091440 free data blocks 64 allocation units of 32768 blocks, 32768 data blocks last allocation unit has 28672 data blocks #</pre>
9	(Optional) After you create the new file system, create a mount point for it. Example: <pre># mkdir /mnt/vg_ntap01_lv1</pre>
10	(Optional) Mount the new file system using the <code>mount</code> command. <code>mount logical_volume mount_point</code> Example: This example mounts the file system you created in the previous step. <pre># mount /dev/vg_ntap01/lvol1 /mnt/vg_ntap01_lv1 # ls -l /mnt/vg_ntap01_lv1 total 0 drwxr-xr-x 2 root root 96 Dec 13 11:40 lost+found #</pre>

Step	Action
11	Verify that the LUN is visible to the host by entering the following command: <pre>sanlun lun show -p all</pre> Example: This example shows the type of output you see using the <code>sanlun</code> command. Note that the output shows the multipathing policy and the PVlinks order of the paths to the LUN. <pre># sanlun lun show -p filer-16:/vol/vol1/lun_04 (LUN 4) VG: /dev/vg_ntap01 36g (38654705664) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/P Multipath-provider: HP PV Links ----- host filer primary partner PVlinks path path /dev/dsk host filer filer path failover state type filename HBA port port priority ----- up primary /dev/dsk/c10t0d4 fcd1 0c 0 up primary /dev/dsk/c4t0d4 fcd0 0c 1 up secondary /dev/dsk/c13t0d4 fcd1 0d 2 up secondary /dev/dsk/c7t0d4 fcd0 0d 3</pre>

Managing LUNs by using VxVM

Prerequisites for using VxVM: To manage LUNs using VxVM, make sure you installed the required VxVM software, Veritas patch bundle, and the ASL. For details, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Putting a new LUN under VxVM management

To put new or additional LUNs under VxVM management, complete the following steps.

Step	Action
1	Make sure you ran the <code>ioscan</code> and <code>ioinit</code> commands to ensure that the host detects the new LUN. See “Getting the host to discover the new LUNs” on page 35 for details.
2	Enter the following command to enable VxVM to detect the new LUN. <pre>vxdctl enable</pre>

Step	Action
3	Enter the following command to verify that VxVM detects the LUN. <pre>vxdisk list</pre> Result: The new disk devices are listed in the output of the <code>vxdisk list</code> command as online invalid disks. If you are using enclosure-based naming, then the NetApp storage model is displayed in the <code>DEVICE</code> column. If you are using disk-based naming, then the controller or disk name (for example, <code>c0t0d3</code>) is displayed. Example: <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:LVM - - LVM Disk_1 auto:LVM - - LVM Filer_0 auto:none - - online invalid Filer_1 auto:none - - online invalid</pre>
4	Enter the following command to initialize the disk for VxVM. <pre>/usr/lib/vxvm/bin/vxdisksetup -i device_name</pre> <i>device_name</i> is the name listed for the LUN in the <code>DEVICE</code> column of the <code>vxdisk list</code> command output you obtained in the previous step. Now when you run <code>vxdisk list</code>, your output should be similar to the following: <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:LVM - - LVM Disk_1 auto:LVM - - LVM Filer_0 auto:none - - online invalid Filer_1 auto:none - - online invalid # vxdisksetup -i FAS30200_0 format=cddisk # vxdisksetup -i FAS30200_1 format=cddisk # vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:LVM - - LVM Disk_1 auto:LVM - - LVM Filer_0 auto:cddisk - - online Filer_1 auto:cddisk - - online</pre>

Step	Action																									
5	Enter the following command to create a disk group for the LUN. <pre>vxdbg init dg_name diskname=device_name</pre> <i>dg_name</i> is the name you assign to the disk group. <i>diskname</i> is the disk name you assign to the device you are adding to the disk group. The <i>diskname</i> represents the device. <i>device_name</i> is the controller name listed for the LUN in the DEVICE column of the vxdisk list command output. Example: <pre># vxdbg init dg1 dg1-disk1=Filer_0</pre>																									
6	Enter the following command to create add disks to the disk group. <pre>vxdbg -g dg_name adddisk diskname=device_name</pre> <i>dg_name</i> is the name of the disk group. <i>diskname</i> is the disk name of the device you are adding to the disk group. The <i>diskname</i> represents the device. <i>device_name</i> is the controller name listed for the LUN in the DEVICE column of the vxdisk list command output. Example: <pre># vxdbg -g dg1 adddisk dg1-disk2=FAS30200_1 # vxdisk list</pre> <table><tr><th>DEVICE</th><th>TYPE</th><th>DISK</th><th>GROUP</th><th>STATUS</th></tr><tr><td>Disk_0</td><td>auto:LVM</td><td>-</td><td>-</td><td>LVM</td></tr><tr><td>Disk_1</td><td>auto:LVM</td><td>-</td><td>-</td><td>LVM</td></tr><tr><td>Filer_0</td><td>auto:cdsdisk</td><td>dg1-disk1</td><td>dg1</td><td>online</td></tr><tr><td>Filer_1</td><td>auto:cdsdisk</td><td>dg1-disk2</td><td>dg1</td><td>online</td></tr></table>	DEVICE	TYPE	DISK	GROUP	STATUS	Disk_0	auto:LVM	-	-	LVM	Disk_1	auto:LVM	-	-	LVM	Filer_0	auto:cdsdisk	dg1-disk1	dg1	online	Filer_1	auto:cdsdisk	dg1-disk2	dg1	online
DEVICE	TYPE	DISK	GROUP	STATUS																						
Disk_0	auto:LVM	-	-	LVM																						
Disk_1	auto:LVM	-	-	LVM																						
Filer_0	auto:cdsdisk	dg1-disk1	dg1	online																						
Filer_1	auto:cdsdisk	dg1-disk2	dg1	online																						

Step	Action
7	Enter the following command to create a logical volume. <pre>vxassist -g <i>dg_name</i> make <i>vol_name</i> <i>size</i></pre> <i>dg_name</i> is the name of the disk group that you defined previously. <i>vol_name</i> is the name you assign to the logical volume. <i>size</i> is the volume size. For example, 50m. Size cannot be equal to or exceed size of the LUN. Example: <pre># vxassist -g dg1 make lun_vol 50m</pre>
8	Create a file system on the volume by entering the following command. <pre>mkfs -F vxfs /dev/vx/rdisk/<i>dg_name</i>/<i>vol_name</i></pre> <i>dg_name</i> is the name of the disk group that you defined previously. <i>vol_name</i> is the name you assigned to the logical volume you defined.
9	Create a mount point for the new file system. Example: <pre># mkdir /mnt/new_lun</pre>
10	Mount the new file system. Example: <pre>mount /dev/vx/dsk/dg1/lun_vol /mnt/new_lun</pre>

Step	Action																																																																																																												
11	Verify the paths to the LUN by entering the following command: <pre>sanlun lun show -p all</pre> Example: The following example shows output for the <code>sanlun</code> command and includes information about the Veritas DMP multipathing policy. If you are using enclosure based naming, the NetApp storage model name is listed as the DMP-NODE. If you are using disk device based naming, the DMP-NODE is listed in device-name format (for example, <code>c16t0d1</code>). Note that the output shows the Veritas path I/O policy as <code>Not Applicable</code>. The policy is single active but <code>sanlun</code> does not identify which path is currently active. <pre># sanlun lun show -p</pre> <pre>fas3020-fak12:/vol/test/fas3020-lun1 (LUN 1) DMP-NODE: (FAS30200_1) 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG Multipath-provider: Veritas DMP</pre> <table><tr><td>host</td><td>filer</td><td></td><td>primary</td><td>partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>/dev/dsk</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>path I/O</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>policy</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c77t0d1</td><td>fcd1</td><td></td><td>0d</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c76t0d1</td><td>fcd1</td><td>0d</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c80t0d1</td><td>fcd0</td><td></td><td>0c</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c79t0d1</td><td>fcd0</td><td>0c</td><td></td></tr></table> <pre>fas3020-fak11:/vol/test/fas3020-lun0 (LUN 0) DMP-NODE: (FAS30200_0) 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG Multipath-provider: Veritas DMP</pre> <table><tr><td>host</td><td>filer</td><td></td><td>primary</td><td>partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>/dev/dsk</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>path I/O</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>policy</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c77t0d0</td><td>fcd1</td><td>0d</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c76t0d0</td><td>fcd1</td><td></td><td>0d</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c80t0d0</td><td>fcd0</td><td>0c</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c79t0d0</td><td>fcd0</td><td></td><td>0c</td></tr></table>	host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	secondary	/dev/dsk/c77t0d1	fcd1		0d	up	primary	/dev/dsk/c76t0d1	fcd1	0d		up	secondary	/dev/dsk/c80t0d1	fcd0		0c	up	primary	/dev/dsk/c79t0d1	fcd0	0c		host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	primary	/dev/dsk/c77t0d0	fcd1	0d		up	secondary	/dev/dsk/c76t0d0	fcd1		0d	up	primary	/dev/dsk/c80t0d0	fcd0	0c		up	secondary	/dev/dsk/c79t0d0	fcd0		0c
host	filer		primary	partner	Veritas																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	secondary	/dev/dsk/c77t0d1	fcd1		0d																																																																																																								
up	primary	/dev/dsk/c76t0d1	fcd1	0d																																																																																																									
up	secondary	/dev/dsk/c80t0d1	fcd0		0c																																																																																																								
up	primary	/dev/dsk/c79t0d1	fcd0	0c																																																																																																									
host	filer		primary	partner	Veritas																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	primary	/dev/dsk/c77t0d0	fcd1	0d																																																																																																									
up	secondary	/dev/dsk/c76t0d0	fcd1		0d																																																																																																								
up	primary	/dev/dsk/c80t0d0	fcd0	0c																																																																																																									
up	secondary	/dev/dsk/c79t0d0	fcd0		0c																																																																																																								

Recommended settings for DMP restore daemon

The values for the DMP restore daemon's policy and interval tuneables determine the type of analysis performed on the paths between the host and storage device, and the frequency of this analysis.

NetApp recommends that you use the following values for these settings:

- ◆ restore daemon interval: 60
- ◆ restore daemon policy: check_all

To set these values, complete the following steps:

Step	Action
1	Log in as root on the host.
2	Display the current values by entering the command: <pre># vxddmpadm stat restored</pre> Example: In this example, the restore daemon interval is set to 300 and the policy is set to check_disabled. <pre># vxddmpadm stat restored The number of daemons running : 1 The interval of daemon: 300 The policy of daemon: check_disabled</pre>
3	Stop the restore daemon by entering the command: <pre># vxddmpadm stop restore</pre>
4	Start the restore daemon with the values set to: ◆ restore daemon interval: 60 ◆ restore daemon policy: check_all Note_____ If you want these new values to persist across reboots, update /sbin/init.d/vxvm-sysboot with desired values. _____ Example: The following command line restarts the restore daemon with the recommended values. <pre>vxddmpadm start restore interval=60 policy=check_all</pre>

Step	Action
5	Verify that the restore daemon is using the correct values. Example: The following command line lets you verify that the restore daemon is using the recommended values. <pre># vxddmpadm stat restored</pre> The number of daemons running : 1 The interval of daemon: 60 The policy of daemon: check_all

Increasing the Powerfail Timeout on disk groups

The VxVM Powerfail Timeout is a timeout value set for a disk group. The default timeout is 30 seconds. If a disk does not respond to requests within 30 seconds, the device is marked as failed. For disk groups containing LUNs, the timeout must be set to 120 seconds so that I/O is not disrupted. To increase the timeout value, complete the following step:

Step	Action
1	For each disk group that contains NetApp LUNs, enter the following command: <pre>/opt/VRTS/bin/vxpfco -g dg_name -t 120</pre>

Displaying LUN and HBA information

Utility for displaying information

The `sanlun` command is a NetApp HP-UX host tool that collects and reports information about host and NetApp storage configuration. For example, you can use the `sanlun` command to show the mapping between LUNs and host devices. You can also use this command to display information about the host HBA.

For detailed information

For detailed information, see the following topics:

- ◆ [“Displaying host LUN information”](#) on page 53
- ◆ [“Displaying host HBA information”](#) on page 57

Displaying LUN and HBA information

Displaying host LUN information

Displaying host LUN information To display host LUN information, complete the following steps.

Step	Action
1	Log in as root on the host.
2	At the host command line, use the <code>sanlun lun show</code> command to display LUN information. This command has two basic formats: ◆ <code>sanlun lun show -p [-v] [filer name:filer pathname filer name all]</code> Note _____ You must use the <code>-p</code> version of the command if you want to see information about the primary and secondary paths. You cannot use the <code>-d</code> option if you use the <code>-p</code> option. ◆ <code>sanlun lun show [-v] [-d host_device_filename all filer_name filer_name:filer_pathname]</code> <code>-p</code> provides information about the primary and secondary paths available to the LUN when you are using multipathing. It lists the multipathing software being used to access the LUN (HP PVlinks or Veritas DMP) and the multipathing policy. For example, if you are using PVlinks, the <code>-p</code> option provides information about path priority. With DMP, the <code>-p</code> option provides information about the DMP multipathing policy. <code>-v</code> option produces verbose output. <code>-d</code> specifies the special device file on the host (<i>host_device_filename</i>); for example, <code>-d /dev/rdsd/c4t0d7</code>. <i>filer_name</i> is the name of the target storage system. <i>filer_pathname</i> is the path to the LUN on the filer. <code>all</code> displays information for all LUNs that are currently attached to the host. If you enter <code>sanlun lun show</code>, <code>sanlun lun show -p</code>, or <code>sanlun lun show -v</code> without any parameters, the utility responds as if you had included the <code>all</code> parameter.

Step	Action
	The command line you enter depends on the information you want to view: ◆ To display a listing of the paths and multipathing policy associated with the LUN, enter the following command line: sanlun lun show -p If you are using the HP-UX LVM, this information from this command line is useful if you need to view path ordering for PVlinks, or to troubleshoot a problem with path ordering. ◆ To display the summary listing of the storage system LUN(s) associated with the host device /dev/rdisk/c3t4d2s2, enter the following command line: sanlun lun show -d /dev/rdisk/c3t4d2s2 ◆ To display verbose output for all the LUNs currently available on the host enter the following command line: sanlun lun show -v all ◆ To display a summary listing of all the LUNs available to the host served by a storage system called filerA, enter the following command line: sanlun lun show filerA ◆ To display a list of paths between the host and a specific LUN on the storage system (for example, lun0) on the storage system filerA, enter the following command line: sanlun lun show filerA:lun_path Note_____ LUNs that were created on the storage system and have not been discovered by the host are not displayed. Also, the device special files must be created. _____ Examples: The following examples show the output you might see on a system using LVM and on a system using DMP. In both cases, keep the following information in mind as you look at the output: ◆ A primary path type indicates a connection that goes to the storage system that owns the LUN. The path can be either via a SAN switch or a direct attached connection. This is the preferred route for I/O to the LUN. A secondary path type indicates a path through a second storage system (i.e., a storage system that does not own the LUN). This path is used only when the primary path to the LUN is unavailable. ◆ The filer port column refers to the port on the filer on which the LUN is accessed. It always represents a primary path. The partner port column refers to the port on the partner in the cluster. It always represents a secondary path. There is also an example showing the output you might see when you use the <code>sanlun</code> command with the <code>-v</code> option.

Step	Action																																																																																																												
	Example using LVM: On a system using LVM, you would see output similar to the following if you executed the <code>sanlun lun show</code> command with the <code>-p</code> option on prior to running the <code>ntap_config_paths</code> utility. In this case, the LUNs are managed by LVM and PVlinks. The PVlinks failover priority column shows the PVlinks order of the paths to the LUN. <pre># sanlun lun show -p filer:/vol/vol1/lun_02 (LUN 2) VG: /dev/vg02 36g (38654705664) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/P Multipath-provider: HP PV Links</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>PVlinks</th></tr><tr><th>path</th><th>path</th><th>/dev/dsk</th><th>host</th><th>filer</th><th>filer</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>path failover</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>priority</th></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c6t0d2</td><td>fcd1</td><td>9a</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c19t0d2</td><td>fcd1</td><td>5a</td><td>not in VG</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c14t0d2</td><td>fcd0</td><td></td><td>9b not in VG</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c12t0d2</td><td>fcd0</td><td></td><td>5b not in VG</td></tr></table> Example using Veritas DMP: On a system using DMP, you would see output similar to the following if you executed the <code>sanlun lun show</code> command. <pre># sanlun lun show -p all filer1:/vol/vol1/lun-9 (LUN 9) DMP-NODE: (c6t1d1) 10g (10737418240) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG Multipath-provider: Veritas DMP</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>Veritas</th></tr><tr><th>path</th><th>path</th><th>/dev/dsk</th><th>host</th><th>filer</th><th>filer</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>path I/O</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>policy</th></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c20t1d1</td><td>td1</td><td></td><td>8b Not Applicable</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c16t1d1</td><td>td1</td><td>8a</td><td>Not Applicable</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c10t1d1</td><td>td0</td><td></td><td>9b Not Applicable</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c6t1d1</td><td>td0</td><td>9a</td><td>Not Applicable</td></tr></table>	host	filer		primary	partner	PVlinks	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path failover						priority	up	primary	/dev/dsk/c6t0d2	fcd1	9a	0	up	primary	/dev/dsk/c19t0d2	fcd1	5a	not in VG	up	secondary	/dev/dsk/c14t0d2	fcd0		9b not in VG	up	secondary	/dev/dsk/c12t0d2	fcd0		5b not in VG	host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	secondary	/dev/dsk/c20t1d1	td1		8b Not Applicable	up	primary	/dev/dsk/c16t1d1	td1	8a	Not Applicable	up	secondary	/dev/dsk/c10t1d1	td0		9b Not Applicable	up	primary	/dev/dsk/c6t1d1	td0	9a	Not Applicable
host	filer		primary	partner	PVlinks																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path failover																																																																																																								
					priority																																																																																																								
up	primary	/dev/dsk/c6t0d2	fcd1	9a	0																																																																																																								
up	primary	/dev/dsk/c19t0d2	fcd1	5a	not in VG																																																																																																								
up	secondary	/dev/dsk/c14t0d2	fcd0		9b not in VG																																																																																																								
up	secondary	/dev/dsk/c12t0d2	fcd0		5b not in VG																																																																																																								
host	filer		primary	partner	Veritas																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	secondary	/dev/dsk/c20t1d1	td1		8b Not Applicable																																																																																																								
up	primary	/dev/dsk/c16t1d1	td1	8a	Not Applicable																																																																																																								
up	secondary	/dev/dsk/c10t1d1	td0		9b Not Applicable																																																																																																								
up	primary	/dev/dsk/c6t1d1	td0	9a	Not Applicable																																																																																																								

Step	Action
	Example using -v: The following example shows an example of output you might see when you enter the <code>sanlun</code> with the <code>-v</code> option only. <pre> # sanlun lun show -v filer: lun-pathname device filename adapter lun size lun state filerX /vol/vol1/3333 /dev/rdisk/c31t0d5 td0 2g (2147483648) GOOD Serial number: OdHNI3miy7QO Filer FCP nodename:50a980000200f2fa Filer FCP portname:50a9800a0200f2cc Filer adapter name: v.4b Filer IP address: 10.97.102.55 Filer volume name:vol1 FSID:0x2500f2cc Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 </pre>

Displaying LUN and HBA information

Displaying host HBA information

Displaying host HBA information To display the host HBA information, complete the following steps.

Step	Action
1	Log in as root on the host.
2	At the host command line, use the <code>sanlun</code> command to display host HBA information. This command has the following format: <pre>sanlun fcp show adapter [-c [-v] [adapter name all]]</pre> <code>-v</code> option produces verbose output. <code>-c</code> gives configuration information that you can use to create igroups. <code>all</code> lists information for all FCP adapters. <i>adapter name</i> lists adapter information for the adapter you specify. Examples: The following example shows the type of output you see when you use different versions of the <code>sanlun fcp show adapter</code> command line. <pre># sanlun fcp show adapter td0 WWPN:50060b000023e2ec td1 WWPN:50060b000023e2cc # sanlun fcp show adapter -c Enter this filer command to create an initiator group for this system: igroup create -f -t hpux "Node-2" 50060b000023e2ec 50060b000023e2cc #</pre>

Step	Action
	<pre># sanlun fcp show adapter -v adapter name: td0 WWPN: 50060b000023e2ec WWNN: 50060b000023e2ed driver name: td model: A6795A model description: Fibre Channel Mass Storage Adapter (PCI) - 2 Gig serial number: Not Available hardware version: 2.3 driver version: @(#) libtd.a HP Fibre Channel Tachyon XL2 Driver B.11.23.0512 \$Date: 2005/09/20 12:22:47 \$Revision: r11.23/1 firmware version: Not Available Number of ports: 1 port type: Fabric port state: Operational supported speed: 2 GBit/sec negotiated speed: 2 GBit/sec OS device name: /dev/td0 adapter name: td1 WWPN: 50060b000023e2cc WWNN: 50060b000023e2cd driver name: td model: A6795A model description: Fibre Channel Mass Storage Adapter (PCI) - 2 Gig serial number: Not Available hardware version: 2.3 driver version: @(#) libtd.a HP Fibre Channel Tachyon XL2 Driver B.11.23.0512 \$Date: 2005/09/20 12:22:47 \$Revision: r11.23/1 firmware version: Not Available Number of ports: 1 port type: Fabric port state: Operational supported speed: 2 GBit/sec negotiated speed: 2 GBit/sec OS device name: /dev/td1</pre>

About this chapter

SAN booting is using a SAN-attached disk, such as a NetApp LUN, as a boot device for a host. SAN booting provides the following advantages:

- ◆ You can remove the hard drives from your servers and use the SAN for your boot needs. This eliminates the costs associated with maintaining and servicing hard drives.
- ◆ Consolidation and centralization—The host uses the SAN for both storage and boot needs.

This chapter provides a section on using SAN booting with LVM and another section on using it with Veritas.

For detailed information about supported components of SAN booting configurations, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. NetApp adds components on an ongoing basis. NetApp qualifies new HBAs, patches and other host components between Host Utilities releases. This online site is updated frequently to provide information about the current stack.

Topics covered in this chapter

This chapter discusses the following topics:

- ◆ “[Overview of SAN booting procedures](#)” on page 60
- ◆ “[\(LVM only\) Creating the boot LUN using LVM](#)” on page 62
- ◆ “[\(LVM only\) Migrating an OS disk using MirrorDisk/UX on PA-RISC systems](#)” on page 68
- ◆ “[\(LVM only\) Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems](#)” on page 75
- ◆ “[\(LVM only\) Testing the alternate boot path for PA-RISC systems](#)” on page 80
- ◆ “[\(LVM only\) Troubleshooting sanboot configurations with PA-RISC systems](#)” on page 81
- ◆ “[\(Veritas only\) Creating the boot LUN using vxcp_lvmroot](#)” on page 83
- ◆ “[\(Veritas only\) Creating the boot LUN using a cold install from the HP-UX DVD](#)” on page 84

Overview of SAN booting procedures

Procedure assumptions

The procedures in this chapter assume the following:

- ◆ You are configuring a PA-RISC server.
Itanium systems use partitions for the boot device. This means that some of the procedures might be different for Itanium systems.
- ◆ The host is booting from an existing internal OS disk. It is not a diskless server.
- ◆ You manage the boot LUN with HP-UX LVM or Veritas VxVM.
- ◆ The SAN Toolkit software is installed on the host.

Note

This chapter contains examples of a few of the procedures you can use to create a bootable SAN LUN. There are other methods for creating a bootable SAN LUN that are not described here that are also supported. For example, you can migrate an existing system disk to a NetApp LUN without the MirrorDisk/UX tool by using `cpio` and `boot device` commands. You can also configure a diskless server to boot from a LUN. For more details, see the documentation for your version of HP-UX.

Procedure overview

To create a bootable SAN LUN on a system, you complete the following general tasks:

- ◆ Create the LUN on which you want to install the HP-UX Operating System (OS). This includes determining the boot LUN's size, mapping it to the host's igroup, and discovering the new LUN on the host.
- ◆ Configure the HP-UX OS on the boot LUN using one of the following methods.

Using LVM:

- ❖ Migrating an existing OS disk to the boot LUN by using MirrorDisk/UX. See [“\(LVM only\) Migrating an OS disk using MirrorDisk/UX on PA-RISC systems”](#) on page 68.
- ❖ Installing the OS by using an Ignite Server or the HP installation CD-ROM. See [“\(LVM only\) Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems”](#) on page 75.

Using Veritas:

- ❖ Creating a Veritas-controlled SAN Boot LUN on HP-UX. See “[\(Veritas only\) Creating the boot LUN using vxcp_lvmroot](#)” on page 83.
- ❖ Performing a cold install directly from the HP-UX Core OS DVD. See “[\(Veritas only\) Creating the boot LUN using a cold install from the HP-UX DVD](#)” on page 84.
- ◆ Set the host’s boot environment by defining the primary and alternate boot paths.

How Itanium systems partition boot devices

For details about creating bootable devices on Itanium systems, see your HP documentation for your OS and server. Basically, Itanium systems partition the boot device into three slices:

- ◆ s0 represents the whole disk, for example c0t0d0s0.
- ◆ s1 represents the EFI (Extensible Firmware Interface) system partition. This is an interface between HP-UX and Itanium BIOS. An example device is c0t0d0s1.
- ◆ s2 represents the HP-UX partition and is equivalent to the whole disk device file for a PA-RISC boot disk. An example device is c0t0d0s2. In a NetApp environment, s2 is used as the device file for the NetApp LUN. This means that if you create a bootable LUN from a NetApp LUN on Itanium systems, you configure the OS on the device file with the s2 designation. In addition, the output of the `ioscan` and `sanlun` commands show the slices of the boot device.

The following `sanlun` output example from an Itanium system using LVM and PVlinks shows the paths to a bootable LUN.

```
filer_a:/vol/vol1/bootlun (LUN 101) VG: /dev/vg00
  18g (19327352832)   lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy: A/P
Multipath-provider: HP PVlinks
```

host	filer		primary	partner	PVlinks	
path	path	/dev/dsk	host	filer	filer	path failover
state	type	filename	HBA	port	port	priority

up	primary	/dev/dsk/c7t12d5s2	td0	4a		0
up	primary	/dev/dsk/c13t12d5s2	td1	5a		1
up	secondary	/dev/dsk/c11t12d5s2	td1		5b	2
up	secondary	/dev/dsk/c5t12d5s2	td0		4b	3

(LVM only) Creating the boot LUN using LVM

- Procedure overview**
- To create a bootable SAN LUN using LVM, you complete the following tasks:
- ◆ Create the LUN on which you want to install the HP-UX Operating System (OS). This includes determining the boot LUN’s size, mapping it to the host’s igroup, and discovering the new LUN on the host.
 - ◆ Configure the HP-UX OS on the boot LUN using one of the following methods:
 - ❖ Migrating an existing OS disk to the boot LUN by using MirrorDisk/UX.
 - ❖ Installing the OS by using an Ignite Server or the HP installation CD-ROM.
 - ◆ Set the host’s boot environment by defining the primary and alternate boot paths.

Determining the size of the boot LUN

If you are installing the OS on the boot LUN: Create a LUN that is large enough to provide enough space for the version of HP-UX you are using. For details, see the documentation for your version of HP-UX.

If you are mirroring an existing host boot disk: Create a LUN that is at least as large as the allocated space used by the existing system boot disk. To determine the size of the system boot disk complete the following steps.

Step	Action
1	On the host, use the <code>vgdisplay</code> command to obtain information about the root volume group (vg00) and determine which disk is used as the system OS disk. <code>vgdisplay -v vg00</code> Result : The output under <code>Physical volumes</code> shows the paths to the system disk. Example: The following example shows the path to the system disk. --- Physical volumes --- PV Name /dev/dsk/clt15d0 PV Status available Total PE 4340 Free PE 2136 Autoswitch On

Step	Action
2	On the host, use the <code>diskinfo</code> command to display the size of the system disk: <pre>diskinfo /dev/rdsk/raw_device_node</pre> <i>raw_device_node</i> is the system disk device in Physical volumes section in the output of the <code>vgdisplay</code> command you used in Step 1. Example: The following example shows that the size of the system disk is approximately 18 GB. Your boot LUN must be at least 18 GB. <pre>diskinfo /dev/rdsk/clt15d0 SCSI describe of /dev/rdsk/clt15d0: vendor: HP 18.2G product id: MAN3184MC type: direct access size: 17783240 Kbytes bytes per sector: 512</pre>

Configuring the NetApp igroup and boot LUN

After you determine the size of the boot LUN, create the boot LUN on the storage system, map it to an igroup, and assign it a LUN ID.

Creating the igroup: If you do not have an igroup set up, create one by following the steps in [“Creating an igroup”](#) on page 24. If you have an existing HP-UX igroup that contains the WWPNs of the host’s HBAs, go to [“Creating the boot LUN”](#) on page 64.

Creating the boot LUN: To create the boot LUN, complete the following steps:

Step	Action
1	On the storage system command line, enter the following command: <pre>lun create -s size -t ostype lun-path</pre> <i>size</i> is the size of the LUN, which you obtained in “Determining the size of the boot LUN” on page 62. You can use one of the following characters to establish the size: c (bytes), w (words or double bytes), b (512), k (kilobytes), m (megabytes), g (gigabytes), t (terabytes). <i>-t ostype</i> indicates the type of operating system running on the host that accesses the boot LUN. Enter hpux. <i>lun-path</i> is the LUN’s path name that includes the volume and qtree. Example: <pre>lun create -s 18g -t hpux /vol/bootvol/bootlun1</pre>
2	Enter the following command to map the LUN to the initiator group. <pre>lun map lun-path igroup-name LUN-ID</pre> <i>lun-path</i> is the path name of the LUN you created. <i>igroup-name</i> is the name of the initiator group you created. <i>LUN-ID</i> is the identification number that you specify. The host associates the LUN ID with the location and path name of the LUN. If you do not specify a LUN ID, Data ONTAP automatically assigns one. The default LUN ID begins with 0 and increments by 1 for each additional LUN that is created. If you have existing LUNs, Data ONTAP assigns the next available LUN ID. Example: <pre>lun map /vol/bootvol/bootlun1 boothost 0</pre>

Step	Action
3	Verify that the LUN is online and mapped to the igroup by entering the following command: <pre>lun show -v</pre> Example: <pre>netapp_1> lun show -v /vol/bootvol/bootlun1 18g (19327352832) (r/w,online, mapped) Serial#: OdCtSYv7Y6w3 Share: none Space Reservation: enabled Multiprotocol Type: hpux Maps: boothost=0</pre>

Discovering the new LUN as a host device

After you create the boot LUN and map it to the igroup, discover the LUN as a host device by completing the following steps:

Step	Action
1	On the host, enter the following command: <pre>ioscan -fn -C disk</pre> Example: In this example, the output from the command shows the following information: - Although the following looks like four different NetApp devices, it is actually just one device with four paths. - There are no device special files for the LUNs yet. - The LUN is visible to the host and S/W state of each path to the LUN is CLAIMED, which means that the path is available. <pre># ioscan -fn -C disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 9 0/1/1/0/1/1.2.0 sdisk CLAIMED DEVICE HP 36.4GATLAS10K3_36_SCA /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0 disk 1 0/4/2/0/4/0.1.8.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 2 0/4/2/0/4/0.1.8.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 3 0/4/2/0/4/0.1.12.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 4 0/4/2/0/4/0.1.12.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 5 0/7/1/0/4/0.1.8.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 6 0/7/1/0/4/0.1.8.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN disk 7 0/7/1/0/4/0.1.12.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN disk 8 0/7/1/0/4/0.1.12.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN</pre>
2	Enter the following command to install special files and create device special files on the host. <pre>ioinit -i</pre> If the <code>ioinit -i</code> command doesn't create device special files, use the <code>insf -e</code> command to install the device special files.

Step	Action	
3	If...	Then...
	You want to migrate an existing system boot disk to the LUN by using MirrorDisk/UX	Proceed to “(LVM only) Migrating an OS disk using MirrorDisk/UX on PA-RISC systems ” on page 68
	You want to install the OS from scratch by using an Ignite Server or an installation CD-ROM	Proceed to “(LVM only) Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems ” on page 75

(LVM only) Migrating an OS disk using MirrorDisk/UX on PA-RISC systems

Mirroring an existing system disk

If you created a new LUN for the boot device, and you want to migrate the existing system disk to the LUN, complete the following steps.

Step	Action																																										
1	On the host, enter the <code>sanlun</code> command to determine which paths to the boot LUN are primary paths. The <code>sanlun</code> command is in the <code>/opt/NetApp/santools/bin</code> directory. You use a primary path from one HBA to the LUN to create a physical volume and add it to the root volume group. sanlun lun show -p Example: In this example, the device <code>c24t0d0</code> is a primary path you can use for the root volume group. This device is used throughout this procedure example as the boot LUN. <pre># sanlun lun show -p boothost:/vol/vol1/bootlun1 (LUN 0) 18g (19327352832) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: Multipath-provider: none</pre> <table><tr><td>host</td><td>filer</td><td></td><td></td><td>primary</td><td>partner</td></tr><tr><td>path</td><td>path</td><td>/dev/dsk</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr></table> <table><tr><td>up</td><td>secondary</td><td>/dev/dsk/c44t0d0</td><td>td1</td><td></td><td>5b</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c46t0d0</td><td>td1</td><td>5a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c26t0d0</td><td>td0</td><td></td><td>4b</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c24t0d0</td><td>td0</td><td>4a</td><td></td></tr></table>	host	filer			primary	partner	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port	up	secondary	/dev/dsk/c44t0d0	td1		5b	up	primary	/dev/dsk/c46t0d0	td1	5a		up	secondary	/dev/dsk/c26t0d0	td0		4b	up	primary	/dev/dsk/c24t0d0	td0	4a	
host	filer			primary	partner																																						
path	path	/dev/dsk	host	filer	filer																																						
state	type	filename	HBA	port	port																																						
up	secondary	/dev/dsk/c44t0d0	td1		5b																																						
up	primary	/dev/dsk/c46t0d0	td1	5a																																							
up	secondary	/dev/dsk/c26t0d0	td0		4b																																						
up	primary	/dev/dsk/c24t0d0	td0	4a																																							
2	Create a bootable physical volume on the LUN by entering the following command: pvcreate -fB /dev/rdsk/ntap_device <i>ntap_device</i> is the raw disk device on the primary path to the LUN. Use the raw device node for the device you found in Step 1. -fB forces the creation of the physical volume and creates space for a boot area on the LUN. Example: <code>pvcreate -fB /dev/rdsk/c24t0d0</code>																																										

Step	Action
3	Add the LUN to the root volume group (vg00). Use a block disk device node, not a raw disk device. <pre>vgextend vg00 /dev/dsk/ntap_device</pre> <i>ntap_device</i> is the disk device file for the primary path to the LUN. Use the device you found in Step 1. Example: <code>vgextend vg00 /dev/dsk/c24t0d0</code>
4	Add the remaining paths to the LUN to the root volume group by using NetApp's <code>ntap_config_paths</code> utility, which is in the <code>/opt/NetApp/santools/bin</code> directory. <pre>ntap_config_paths vg00</pre>
5	Enter the following command to copy boot utilities to the boot LUN. <pre>mkboot /dev/rdisk/ntap_device</pre> <i>ntap_device</i> is the device used for the boot LUN, for example, <code>c24t0d0</code>.
6	Enter the following commands to make sure that the host will boot if either the boot LUN or the system disk is unavailable. <pre>mkboot -a "hpux -lq" /dev/dsk/system_disk_device mkboot -a "hpux -lq" /dev/dsk/ntap_device</pre> <i>system_disk_device</i> is the path to the system disk device node, for example, <code>c1t15d0</code>. <i>ntap_device</i> is the device node on the primary path to the boot LUN, for example <code>c24t0d0</code>.

Step	Action
7	Enter the following command to display the logical volumes on the existing system boot disk (vg00). The logical volumes are displayed under the --- Logical volumes --- line in the command output. You mirror the logical volumes to the boot LUN. <pre>vgdisplay -v vg00</pre> Example: The following example is partial output from the <code>vgdisplay</code> command to show how two logical volumes are listed. <pre>--- Logical volumes --- LV Name /dev/vg00/lvol1 LV Status available/syncd LV Size (Mbytes) 300 Current LE 75 Allocated PE 75 Used PV 1 LV Name /dev/vg00/lvol2 LV Status available/syncd LV Size (Mbytes) 2048 Current LE 512 Allocated PE 512 Used PV 1</pre>
8	For each logical volume, use the <code>lvextend</code> command to mirror the logical volume from the existing system boot disk (vg00) to the boot LUN. <pre>lvextend -m 1 lv_path /dev/dsk/ntap_device</pre> <code>-m 1</code> specifies to mirror one logical volume. <i>lv_path</i> is the path to the logical volume, for example, <code>/dev/vg00/lvol2</code> <i>ntap_device</i> is the device on the primary path to the boot LUN, for example, <code>c24t0d0</code>. Example: This example mirrors logical volume 1 to the boot LUN (<code>/dev/dsk/c24t0d0</code>). <pre>lvextend -m 1 /dev/vg00/lvol1 /dev/dsk/c24t0d0</pre> The newly allocated mirrors are now being synchronized. This operation will take some time. Please wait Logical volume <code>"/dev/vg00/lvol1"</code> has been successfully extended. Volume Group configuration for <code>/dev/vg00</code> has been saved in <code>/etc/lvmconf/vg00.conf</code>
9	Update the Boot Device Reserved Area (BDRA) by entering the following command: <pre>lvlnboot -R</pre>

Step	Action
10	Using <code>vi</code> or another text editor, add a line to the <code>/stand/bootconf</code> file to identify the boot disk to the SDUX utility. Use the path to the block device node. <pre>1 /dev/dsk/<i>ntap_device</i></pre> <i>ntap_device</i> is the device used for the boot LUN. Example: 1 /dev/dsk/c24t0d0

Setting the host's boot paths

After you mirror the boot disk to the boot LUN, set the primary and alternate boot paths for the host by using the `setboot` command on the host. When you use the `setboot` command, you specify hardware paths to a device. This means that you have to identify the primary and secondary paths to the boot LUN by using the `sanlun` command, and then identify which hardware paths they correspond to by using the `ioscan` command.

To ensure system availability, the primary boot path must use a primary path to the boot LUN. The alternate boot path must use a secondary path to the boot LUN.

When you define the primary boot path, make sure you use the device that you identified in [Step 1](#) of “[Mirroring an existing system disk](#)” on page 68. This is the device you added to the root volume group and copied the HP-UX boot utilities to.

To set the host boot paths, complete the following steps.

Step	Action																																																
1	Display the paths to the boot LUN by entering the following command: <pre>sanlun lun show -p</pre> Example using LVM: In this example from a host using LVM, the primary path to the boot LUN from host is /dev/dsk/c24t0d0. This is the device to which you mirrored logical volumes and copied HP-UX boot utilities in “Mirroring an existing system disk” on page 68. A secondary path to the boot LUN from the host is /dev/dsk/c44t0d0. These paths are used for the remainder of this procedure example. <pre># sanlun lun show -p boothost:/vol/vol1/bootlun (LUN 0) VG: /dev/vg00 18g (19327352832) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/P Multipath-provider: HP PVlinks</pre> <table><tr><td>host</td><td>filer</td><td></td><td></td><td>primary</td><td>partner</td></tr><tr><td>path</td><td>path</td><td>/dev/dsk</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td colspan="6">-----</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c24t0d0</td><td>td0</td><td>4a</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c46t0d0</td><td>td1</td><td>5a</td><td>1</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c44t0d0</td><td>td1</td><td>5b</td><td>2</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c26t0d0</td><td>td0</td><td>4b</td><td>3</td></tr></table>	host	filer			primary	partner	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port	-----						up	primary	/dev/dsk/c24t0d0	td0	4a	0	up	primary	/dev/dsk/c46t0d0	td1	5a	1	up	secondary	/dev/dsk/c44t0d0	td1	5b	2	up	secondary	/dev/dsk/c26t0d0	td0	4b	3
host	filer			primary	partner																																												
path	path	/dev/dsk	host	filer	filer																																												
state	type	filename	HBA	port	port																																												

up	primary	/dev/dsk/c24t0d0	td0	4a	0																																												
up	primary	/dev/dsk/c46t0d0	td1	5a	1																																												
up	secondary	/dev/dsk/c44t0d0	td1	5b	2																																												
up	secondary	/dev/dsk/c26t0d0	td0	4b	3																																												

Step	Action
2	Enter the following command to determine which hardware paths correspond to the primary and secondary paths to the boot LUN, which you displayed in step 1. <pre>ioscan -funC disk</pre> Example: In this example, the hardware path for the primary path to the boot LUN is 0/4/0/0.1.4.0.0.0.0. Use this hardware path for the host's primary boot path. The hardware path for the secondary path to the boot LUN is 0/6/2/0.1.7.0.0.0.0. Use this hardware path for the host's alternate boot path. <pre># ioscan -funC disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC /dev/dsk/c1t15d0 /dev/rdisk/c1t15d0 disk 14 0/4/0/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c24t0d0 /dev/rdisk/c24t0d0 disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c26t0d0 /dev/rdisk/c26t0d0 disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c46t0d0 /dev/rdisk/c46t0d0 disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c44t0d0 /dev/rdisk/c44t0d0</pre>
3	Set the host's primary boot path by entering the following command: <pre>setboot -p prim_hw_path_to_lun</pre> <i>prim_hw_path_to_lun</i> is the primary hardware path to the boot LUN from the host's HBA1, which you displayed in step 2. Example: The following example shows how to set the primary boot path to the boot LUN by using the hardware path that corresponds to /dev/dsk/c24t0d0, the primary path. <pre>setboot -p 0/4/0/0.1.4.0.0.0.0</pre>

Step	Action	
4	Set the host's alternate boot path by entering the following command: <pre>setboot -a <i>sec_hw_path_to_lun</i></pre> <i>sec_hw_path_to_lun</i> is the secondary hardware path to the boot LUN from the host's HBA2, which you displayed in step 2. Example: The following example shows how to set the alternate boot path to the boot LUN by using the hardware path that corresponds to /dev/dsk/c44t0d0, the secondary path. <pre>setboot -a 0/6/2/0.1.7.0.0.0.0</pre>	
5	Test the primary boot path by shutting down and rebooting the system. Note_____ If Autoboot is disabled (set to OFF) for your system, you will have to boot the system from the Boot Control Handler (BCH) menu. _____ <pre>shutdown -r -y 0</pre>	
6	If...	Then...
	You cannot reboot the system.	Proceed to “(LVM only) Troubleshooting sanboot configurations with PA-RISC systems” on page 81
	You can reboot the system	Proceed to “(LVM only) Testing the alternate boot path for PA-RISC systems” on page 80.

(LVM only) Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems

Identifying paths to the boot LUN When you install the OS on the boot LUN, you need to specify a primary path to the boot LUN as the Root Disk. The installation process uses the `setboot` command to define the primary path to the boot LUN as the host's primary boot path. After the installation is complete, you manually enter the `setboot` command to define an alternate boot path. This means that you have to identify the primary and secondary paths to the boot LUN, and their corresponding hardware paths, before you begin the installation.

Note You use the same procedure for PA-RISC systems and for Itanium systems.

To identify the primary and secondary paths to the boot LUN, complete the following steps.

Step	Action
1	Display the primary and secondary paths to the boot LUN by entering the <code>sanlun</code> command, which is in the <code>/opt/NetApp/santools/bin</code> directory: <code>sanlun lun show -p</code> Example: In this example, a primary path to the boot LUN from the host is <code>/dev/dsk/c24t0d0</code>. A secondary path to the boot LUN from the host is <code>/dev/dsk/c44t0d0</code>. These paths are used as primary and secondary (alternate) paths for the remainder of this procedure example. # sanlun lun show -p boothost:/vol/vol1/bootlun (LUN 0) 18g (19327352832) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: Multipath-provider: ----- host filer primary partner path path /dev/dsk host filer filer state type filename HBA port port ----- up primary /dev/dsk/c24t0d0 td0 4a up primary /dev/dsk/c46t0d0 td1 5a up secondary /dev/dsk/c44t0d0 td1 5b up secondary /dev/dsk/c26t0d0 td0 4b

Step	Action
2	Enter the following command to determine which hardware paths correspond to the primary and secondary paths to the boot LUN, which you displayed in Step 1. <pre>ioscan -funC disk</pre> Example: In this example, the hardware path for the primary path to the boot LUN is 0/4/0/0.1.4.0.0.0.0. Use this hardware path for the host's primary boot path. The hardware path for the secondary path to the boot LUN is 0/6/2/0.1.7.0.0.0.0. Use this hardware path for the host's alternate boot path. <pre># ioscan -funC disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC /dev/dsk/c1t15d0 /dev/rdisk/c1t15d0 disk 14 0/4/0/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c24t0d0 /dev/rdisk/c24t0d0 disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c26t0d0 /dev/rdisk/c26t0d0 disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c46t0d0 /dev/rdisk/c46t0d0 disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c44t0d0 /dev/rdisk/c44t0d0</pre>
3	Record the primary and secondary hardware and device paths. During installation, you use a primary path to the LUN to define the root disk (vg00). After you install the OS, you use a secondary path to define the host's alternate boot path.

Installing the OS on the boot LUN

To install the OS on the boot LUN, complete the following steps.

Note

If you are using an Itanium system, see the HP documentation for instructions.

1	If...	Then...
	You are installing from a CD	Follow the HP instructions to boot from the CD.
	You are installing from an Ignite Server	1. Shutdown the host by entering the following command: shutdown -r -y 0 Note _____ If Autoboot is enabled (set to ON), you have to interrupt the boot process to get to the Boot Control Handler (BCH) menu. If Autoboot is disabled (set to OFF), the system comes up and displays the BCH menu. _____ 2. When the system is back up, enter the following command at the BCH Main Menu prompt: bo lan.<server_IP> <i>server_IP</i> is the IP address of the Ignite Server. Example: Main Menu: Enter command or menu > bo lan.10.105.36.44 Interact with IPL (Y, N, or Cancel)?> n

2	In the HP-UX installation, configure system parameters, such as network services, file-system characteristics, and other configuration details. Make sure that you define the Root Disk as the primary path to the boot LUN. This is the path you obtained in Step 1 of “Identifying paths to the boot LUN” on page 75. In the Ignite-UX Interface, you define the Root Disk in the Basic tab. When you define the primary path to the boot LUN as the Root Disk, the installation script does the following: ◆ Configures the root volume group (vg00) and adds remaining path to the LUN to the root volume group. ◆ Uses the <code>setboot -p</code> command to define the primary path to the boot LUN as the host’s primary boot path. Note_____ For details about configuring other system parameters, consult the HP documentation for your OS. _____
3	After you complete the HP-UX installation, install the FCP HP-UX Host Utilities software on the boot LUN.
4	Add the remaining paths to the LUN to the root volume group by using NetApp’s <code>ntap_config_paths</code> command, which is in the <code>/opt/NetApp/santools/bin</code> directory. ntap_config_paths vg00
5	Set the host’s alternate boot path by entering the following command: setboot -a <i>sec_hw_path_to_lun</i> <i>sec_hw_path_to_lun</i> is the secondary hardware path to the boot LUN from the host, which you obtained in Step 2 of “Identifying paths to the boot LUN” on page 75. Example: <code>setboot -a 0/6/2/0.1.7.0.0.0.0</code>
6	Test the primary boot path by shutting down and rebooting the system. Note_____ If Autoboot is disabled (set to OFF) for your system, you will have to boot the system manually. _____ shutdown -r -y 0

7	If...	Then...
	You cannot reboot the system.	Proceed to “(LVM only) Troubleshooting sanboot configurations with PA-RISC systems” on page 81
	You can reboot the system	Proceed to “(LVM only) Testing the alternate boot path for PA-RISC systems” on page 80.

(LVM only) Testing the alternate boot path for PA-RISC systems

Testing the alternate boot path

NetApp recommends that you test the alternate boot path before you put a SAN-booted configuration into a production environment. To test the alternate boot path, complete the following steps.

Step	Action
1	Shut down the host by entering the following command: <code>shutdown -r -y 0</code>
2	At the Main Menu prompt of the Boot Console Handler (BCH), boot from the alternate boot path by entering the following command: <code>bo alt</code>
3	Now stop the alternate boot path and go back to the primary boot patch. Shut down the host again: <code>shutdown -r -y 0</code>
4	At the Main Menu prompt of Boot Console Handler (BCH), boot from the primary boot path by entering the following command: <code>bo pri</code>

(LVM only) Troubleshooting sanboot configurations with PA-RISC systems

If you cannot boot from the primary or alternate path

If the system does not boot from the primary or alternate path, check the following configuration components:

- ◆ Connectivity—Check physical connectivity and cabling between the HBA, the switches, and the NetApp cluster.
- ◆ Boot path settings—If your system fails to boot from either the primary or alternate path, boot from the internal system disk, and then check your boot path settings.

Checking boot path settings: To boot from the system disk and check your boot path settings, complete the following steps:

Step	Action
1	If the system does not return to the Boot Control Handler (BCH) menu, reset the system to return to the BCH prompt.
2	Boot from the internal system disk by entering the following command: bo <i>path</i> <i>path</i> is the hardware path to the internal disk that the system formerly booted from. Example: The following example uses the hardware path to the internal system disk. bo 0/0/1/1.15.0
3	Verify that you correctly translated the primary and secondary device paths to the boot LUN to the correct hardware paths by using the following commands: ◆ <code>sanlun lun show -p</code>—Shows the primary and secondary device paths to the boot LUN. ◆ <code>ioscan -funC disk</code>—Shows the hardware paths to the device files.

Step	Action
4	Make sure that all paths to the boot LUN are in the root volume group by entering the following command: <pre>vgdisplay -v vg00</pre> Result: Paths to the LUN are listed under the Physical volumes section of the output. Example: In the following example, all paths to the boot LUN are in the root volume group. <pre>--- Physical volumes --- PV Name /dev/dsk/c24t0d0 PV Name /dev/dsk/c26t0d0 Alternate Link PV Name /dev/dsk/c46t0d0 Alternate Link PV Name /dev/dsk/c44t0d0 Alternate Link PV Status available Total PE 4606 Free PE 1577 Autoswitch On</pre>
5	If all paths to the boot LUN are not in the root volume group, add the remaining paths by using the <code>ntap_config_paths</code> utility, which is in the <code>/opt/NetApp/santools/bin</code> directory. <pre>ntap_config_paths vg00</pre>

(Veritas only) Creating the boot LUN using vxcp_lvmroot

Prerequisites specific to Veritas

Before you begin this procedure may sure you have the following:

- ◆ Veritas and the NetApp Array Support Library (ASL) should already be installed on the host. See “[Veritas and Array Support Library](#)” on page 115 for information on installing ASL.
- ◆ The NetApp LUN that you plan to use for the new SAN boot device should already exist on the NetApp storage system and be visible to the host. See “[Creating, Configuring, and Managing LUNs](#)” on page 21 for information on creating and working with LUNs.
- ◆ The NetApp LUN should have been discovered and be usable by the HP-UX operating system. See “[Creating, Configuring, and Managing LUNs](#)” on page 21 for more information.

Using vxcp_lvmroot to create the boot LUN

You can use `vxcp_lvmroot` to create a Veritas-controlled SAN boot LUN on HP-UX. Complete the following steps:

Step	Action
1	Copy the existing LVM root disk onto the NetApp LUN by following the <code>vxcp_lvmroot</code> procedure contained in the Veritas Volume Manager 4.1 Administrator's Guide for HP-UX.
2	Verify that the <code>vxpfto</code> and the restore daemon settings are tuned according to the NetApp recommendations. See “ Increasing the Powerfail Timeout on disk groups ” on page 51 and “ Recommended settings for DMP restore daemon ” on page 50 for more information.
3	Configure the host's primary and alternate boot path settings using the <code>setboot</code> command

(Veritas only) Creating the boot LUN using a cold install from the HP-UX DVD

Prerequisite specific to Veritas and creating a boot LUN using a cold install

Before you begin this procedure make sure that the NetApp LUN that you plan to use for the new SAN boot device already exists on the NetApp storage system and is visible to the host. See “[Creating, Configuring, and Managing LUNs](#)” on page 21 for information on creating and working with LUNs.

Caution _____
The HP-UX host **must** see only one path to the LUN or the installation will fail. You can use zoning to accomplish this.

Performing a cold install to create the boot LUN

You can create a Veritas-controlled SAN boot LUN by performing a cold install directly from the HP-UX Core OS DVD. Complete the following steps:

Step	Action
1	Perform an advanced Ignite-UX installation to the NetApp LUN. Caution _____ Make sure that, when you select the root disk, Ignite-UX displays only one path to the NetApp LUN.
2	After the Ignite-UX installation completes, install the Veritas Storage Foundation on the host. There should still only be one path the NetApp LUN.
3	Install the NetApp ASL on the host. For more information, see “ Veritas and Array Support Library ” on page 115.
4	Install the SAN Toolkit on the host. For more information, see “ Installing the Host Utilities ” on page 11.
5	Make the other paths to the NetApp LUN visible to the host.
6	Run the appropriate HP-UX and VxVM commands to discover the and use the new paths. You can use <code>ioscan</code> , <code>infosf</code> , and <code>vxdctl enable</code> .

Step	Action
7	Verify that the <code>vxpfto</code> and restore daemon settings are tuned according to the NetApp recommendations. See “ Increasing the Powerfail Timeout on disk groups ” on page 51 and “ Recommended settings for DMP restore daemon ” on page 50 for more information.
8	Configure the host's primary and alternate boot path settings using the <code>setboot</code> command.

About this chapter

This chapter describes how to troubleshoot software installation problems and HBA and LUN management problems.

Topics covered in this chapter

This chapter discusses the following topics:

- ◆ [“Troubleshooting checklists”](#) on page 88
- ◆ [“Correcting problems with Veritas”](#) on page 91
- ◆ [“Correcting volume group multipathing with ntap_config_paths”](#) on page 92
- ◆ [“Handling a takeover and giveback”](#) on page 94
- ◆ [“Recovering from a host reboot during a takeover.”](#) on page 95
- ◆ [“Enabling the host to identify more than 8 LUNs per target HBA”](#) on page 98
- ◆ [“Commands you can use to troubleshoot problems”](#) on page 104
- ◆ [“NetApp’s diagnostic utilities”](#) on page 109
- ◆ [“General troubleshooting”](#) on page 114

Troubleshooting checklists

Connectivity problems when using LVM

On HP-UX systems using LVM, applications don't generally produce I/O errors. Instead, I/O retries continue uninterrupted if connectivity is lost. If you have a situation where I/O retries appear to be continuing without succeeding, check the following areas for connectivity problems:

- ◆ NetApp system setup
- ◆ Connectivity between the NetApp system and the Fibre Channel switch
- ◆ Fibre Channel switch setup
- ◆ Fibre Channel connectivity to the host
- ◆ Host setup

NetApp storage system setup checklist

Make sure that the following situations are true on your storage system:

If you have a cluster, make sure the FCP cfmodes setting is correct and supported by your configuration. To determine which cfmodes are supported, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Note

Changing the cfmodes affects the volume group and zoning on your system. If you need to change the volume group, contact customer support.

- ◆ FCP is running and all the target ports are in a good state.

Note

Keep in mind that if an HP-UX FCP driver is not in use, it logs off the target port. For this reason, the `fcv show initiators` command might not display all the host drivers as connected even though they are.

- ◆ The igroups have the following values for the attributes:
 - ❖ The `ostype` attribute must be `hpux`
 - ❖ The `vsa` attribute must be set to `yes`

Note

This attribute only appears in the output displays if it is incorrectly set. When it is correctly set, no references to it appear.

- ◆ For the LUNs, the following is true:
 - ❖ The `ostype` attribute is `hpux`
 - ❖ They are mapped to host igroups.
 - ❖ They are in a good state (that is, when you execute the `sanlun lun show -p` command, it lists the LUN state as “good”)

Host setup checklist

The following is a checklist of items that pertain to the host setup. You should check to see that the following is true:

- ◆ The HBAs are working.
Check the output from the `sanlun fcp show adapter -v` command.
- ◆ The host is seeing all of the LUNs it should.
- ◆ The software states reported by the `ioscan` command for the LUNs are “CLAIMED” (that is, in a good state).
- ◆ The `sanlun lun show -p` command is not reporting problems such as the following:
 - ❖ Downed paths
 - ❖ Missing paths
 - ❖ Out of order paths
 - ❖ The NetApp cluster is in a takeover state
- ◆ There are no suspicious messages in the `/var/adm/syslog/syslog.log` file, including messages about the following:
 - ❖ I/O timeouts
 - ❖ Reported disk or SCSI errors
 - ❖ EMS event notifications

Note

A suspicious message might simply say that there is a loss of connectivity without explaining what the problem is.

- ◆ There are no other volume group configuration problems reported.
Use the `vgdisplay -v` command to display information about volume group configurations.

Situations affecting volume groups and disk devices

If certain events or changes occur, you must reconfigure the volume groups on the HP-UX host and disk devices the host sees from the storage system. These situations include:

- ◆ If the host reboots while a NetApp cluster is in a degraded or takeover state (see [“Recovering from a host reboot during a takeover.”](#) on page 95)
- ◆ If you change any of the following settings on the storage system:
 - ❖ FCP cfmode (do not do this unless directed to do so by customer support)
 - ❖ igroup VSA setting

If you are using LVM, see [“Correcting volume group multipathing with ntap_config_paths”](#) on page 92.

Fixing problems with the volume groups online

You can resolve problems with a volume group while it is online as long as you have at least one active path to each physical volume in the volume group. For example, if there is a missing path, you can locate it without taking the volume group offline.

Fixing problems with the volume groups offline

If all the paths to a physical volume go down and new paths appear, you must perform a host volume group reconfiguration. The following situations require a host volume group reconfiguration:

- ◆ The cfmode needs to change. (Do not change the cfmode unless directed to do so by NetApp.)
- ◆ The `vsa` attribute for the igroup is not enabled. This would happen if you created an igroup without specifying `hpux` as the operating system type. By default, this value is set correctly when you create an HP-UX igroup. (For information about setting the `vsa` attribute, see [“Changing the vsa value”](#) on page 98.)

Correcting problems with Veritas

Application I/O failures occur when using Veritas

In some cases you may see application I/O failures occurring during storage system cluster operations, such as takeover or giveback.

If this happens, perform the following checks:

- ◆ Verify that the `vxpfto` setting is set to the values specified in this document. See “[Increasing the Powerfail Timeout on disk groups](#)” on page 51.
- ◆ Verify that the DMP restore daemon settings are set to the values specified in this document. See “[Recommended settings for DMP restore daemon](#)” on page 50.
- ◆ Verify the storage system is configured with a supported ONTAP cfmode. See the *Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Enclosure-based naming is not working correctly when using Veritas

You are using enclosure-based naming on your system, and it is not working correctly.

Note

Using enclosure-based naming is optional. The Host Utilities do not require it.

If this happens, verify that:

- ◆ The ASL is loaded correctly. See “[Veritas and Array Support Library](#)” on page 115.
- ◆ You have changed the disk naming scheme to enclosure-based naming using the `vxdiskadm` utility

Correcting volume group multipathing with ntap_config_paths

Using the utility ntap_config_paths with LVM

When you’re using HP LVM with PVlinks, you can use the `ntap_config_paths` utility that NetApp provides to perform the following tasks:

- ◆ Set up alternate paths to an LVM volume group after you create the volume group with one primary path.
- ◆ Correct problems with path order within a volume group.

Example: The following example shows how `ntap_config_paths` works to correct paths to volume groups. The `sanlun lun show -p` command output includes the comment “not in VG,” which means that the LUN is a member of a volume group, but some of the paths to it are missing.

Note This is an abbreviated example to make it easier to read. Part of the `ntap_config_paths` output was removed.

```
# sanlun lun show -p all

filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
      2g (2147483648))          lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy: A/P
Multipath-provider: HP PVlinks

-----
host      filer                                primary partner PVlinks
path      path                                /dev/dsk      host filer  filer  failover
state     type                                filename      HBA  port   port   priority
-----
up        primary                            /dev/dsk/c34t0d5  td0  4a          0
up        primary                            /dev/dsk/c35t0d5  td1  5a          not in VG
up        secondary                          /dev/dsk/c30t0d5  td1          5b  not in VG
up        secondary                          /dsk/dsk/c31t0d5  td0          4b  not in VG

# ntap_config_paths

Getting information from sanlun...

Adding missing path with:
    vgextend /dev/vg_ntap01 /dev/dsk/c30t0d5
(output removed)
Moving VG path to the end of the alternate paths with:
    vgreduce /dev/vg_ntap01 /dev/dsk/c34t0d5
```

```
vgextend /dev/vg_ntap01 /dev/dsk/c34t0d5  
(output removed)  
Moving VG path to the end of the alternate paths with:  
vgreduce /dev/vg_ntap01 /dev/dsk/c31t0d5  
vgextend /dev/vg_ntap01 /dev/dsk/c31t0d5  
(output removed)
```

Handling a takeover and giveback

Handling a takeover and giveback

If a storage system fails, a takeover occurs. The host can no longer use its primary paths to access the LUNs owned by the failed storage system.

If you have a storage system cluster, the host can use the secondary paths to access the LUNs on the failed storage system. These secondary paths go through the storage system that did not fail.

Note

If the host reboots during the takeover, it loses the primary paths to LUNs. This is because the storage system loses its optimized paths to the LUNs. If a reboot occurs, see the section “[Recovering from a host reboot during a takeover](#)” on page 95 to determine you which steps you must perform to get that storage system to rediscover the optimized paths.

What happens during a takeover

Several things happen during a takeover:

- ◆ All the primary paths from the host that go through the failed storage system to its LUNs become unresponsive. If you execute the `ioscan` command, it displays a SW State of `NO_HW`, which means unresponsive paths.
- ◆ Applications do not see I/O errors from the system.
- ◆ It might take PVlinks up to $(P * T)$ seconds to resume I/O on the secondary paths during a takeover, where
 - ❖ P is the number of primary paths.
 - ❖ T is the I/O timeout on the physical volume. To see the value, enter `pvdisplay -v`; to change this value, enter `pvchange`.
- ◆ There might be timeout messages in the host’s syslog file.

What happens during a giveback

When the giveback occurs, the host automatically starts accessing the LUNs using the primary paths.

Recovering from a host reboot during a takeover.

Events that occur if a host reboots during a takeover

- If a host reboots while the storage system is in a takeover state, the following events occur:
- 1. The downed storage system loses the primary paths to the LUNs it owns. It may lose its secondary paths as well.
 - 2. The volume groups deactivate those primary paths.
 - 3. The volume groups remain active, but are aware of the secondary paths only. These are the paths from the active storage system to the LUNs on the failed storage system.
 - 4. When the giveback occurs, the host is only aware of the secondary paths; it uses the proxy paths, not the primary paths.

When a reboot occurs, you need to perform steps so that the host rediscovers the primary paths to the volume groups.

If no host reboot occurs during the takeover, the host automatically notices when the primary paths come back up and transfers information back to them without any user intervention. See “[Handling a takeover and giveback](#)” on page 94 for more information.

Recovering from a host reboot during a takeover

If the host reboot occurs while the cluster is in a takeover state, it does not automatically rediscover the paths to the LUNs after giveback. After the giveback occurs, perform the following steps for each volume group hosted on the storage system.

Step	Action
1	Repair the problem that caused the filer to fail, and then reboot the storage system.
2	On the NetApp console, perform the giveback by entering the following command. cf giveback

Step	Action																																										
3	On the host, enter the <code>ioscan</code> command to detect the primary paths. Example: The following is an example of the output from the <code>ioscan</code> command that shows all paths to the LUN. <pre># ioscan -funC disk Class I H/W Path Driver S/W State H/W Type Description ===== disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC /dev/dsk/c1t15d0 /dev/rdsk/c1t15d0 disk 14 0/4/0/0.1.4.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c24t0d0 /dev/rdsk/c24t0d0 disk 28 0/4/0/0.1.7.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c26t0d0 /dev/rdsk/c26t0d0 disk 37 0/6/2/0.1.4.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c46t0d0 /dev/rdsk/c46t0d0 disk 46 0/6/2/0.1.7.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN /dev/dsk/c44t0d0 /dev/rdsk/c44t0d0</pre>																																										
4	On the host, enter the following command to confirm that the host detects all paths to the LUNs. sanlun lun show -p all Result: The <code>sanlun</code> command shows all recovered paths, but the paths are not shown in the correct multipathing order. Example: The following shows partial <code>sanlun</code> command output after you have used the <code>ioscan</code> command to detect the recovered paths. Note that the PVlinks priority is not correct—a secondary path is listed as the first path in the PVlinks order. <pre># sanlun lun show -p filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01 2g (2147483648)) lun state: GOOD Filer_CF_State: Cluster Enabled Array_Policy: A/P Multipath-provider: HP PVlinks</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>PVlinks</th></tr><tr><th>path</th><th>path</th><th>/dev/dsk host</th><th>filer</th><th>filer</th><th>failover</th></tr><tr><th>state</th><th>type</th><th>filename HBA</th><th>port</th><th>port</th><th>priority</th></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c44t0d0</td><td>td1</td><td>5b</td><td>0 (Wrong)</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c24t0d0</td><td>td0</td><td>4a</td><td>1</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c46t0d0</td><td>td1</td><td>5a</td><td>2</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c26t0d0</td><td>td0</td><td>4b</td><td>3</td></tr></table>	host	filer		primary	partner	PVlinks	path	path	/dev/dsk host	filer	filer	failover	state	type	filename HBA	port	port	priority	up	secondary	/dev/dsk/c44t0d0	td1	5b	0 (Wrong)	up	primary	/dev/dsk/c24t0d0	td0	4a	1	up	primary	/dev/dsk/c46t0d0	td1	5a	2	up	secondary	/dev/dsk/c26t0d0	td0	4b	3
host	filer		primary	partner	PVlinks																																						
path	path	/dev/dsk host	filer	filer	failover																																						
state	type	filename HBA	port	port	priority																																						
up	secondary	/dev/dsk/c44t0d0	td1	5b	0 (Wrong)																																						
up	primary	/dev/dsk/c24t0d0	td0	4a	1																																						
up	primary	/dev/dsk/c46t0d0	td1	5a	2																																						
up	secondary	/dev/dsk/c26t0d0	td0	4b	3																																						

Step	Action
5	Enter the following command to reconfigure the PVlinks path order. ntap_config_paths
6	Enter the <code>ntap_config_paths</code> command a second time. You must run this command again because in this situation it does not set the correct path order when you enter it the first time. ntap_config_paths
7	Enter the <code>sanlun</code> command to verify that the paths are configured in the correct order. sanlun lun show -p all Example: The following partial output shows the correct order of the paths to the LUN. For example, the primary paths to the LUN are the first two paths listed. <pre># sanlun lun show -p all filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01 2g (2147483648)) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/P Multipath-provider: HP PVlinks ----- host filer primary partner PVlinks path path /dev/dsk host filer filer path failover state type filename HBA port port priority ----- up primary /dev/dsk/c24t0d0 td0 4a 0 up primary /dev/dsk/c46t0d0 td1 5a 1 up secondary /dev/dsk/c44t0d0 td1 5b 2 up secondary /dev/dsk/c26t0d0 td0 4b 3</pre>

Enabling the host to identify more than 8 LUNs per target HBA

Using vsa to change the number of LUNs identified

The igroup `vsa` attribute enables the HP-UX host to identify more than eight LUNs per target HBA.

By default, this attribute is set when you create an igroup that specifies `hpux` as the operating system type. You must manually set this attribute if you change an existing igroup to an HP-UX igroup by modifying the `ostype` attribute.

Note_____

You should never turn this attribute off.

Changing the vsa value

To manually set the `vsa` attribute to on, complete the following steps. You will need to perform some of these on the host and some on the storage system.

Step	Action
1	(Host) 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. You should record the LUN volume group membership. To find all the LUNs and volume groups that contain LUNs from a specific pair of NetApp storage systems, run the <code>sanlun lun show -p</code> command for each filer: <pre>sanlun lun show -p <filer1 name> sanlun lun show -p <filer2 name></pre>
2	(Host) Close the logical volumes. If a logical volume contains a file system, unmount the file system.

Step	Action
3	(Host) Deactivate each volume group by entering the following command: <pre>vgchange -a n /dev/vol_group_name</pre> <i>vg_name</i> is the path to the volume group. Example: The following command deactivates the volume group /dev/ntap01: <pre># vgchange -a n /dev/ntap01</pre>
4	(Host) Remove the volume group entry from the /etc/lvmtab and associated device files by using the vgexport command: <pre>vgexport /dev/ntap01</pre>
5	(Storage system) Go into diagnostic mode by entering the following command: <pre>priv set diag</pre> Example: When you start diagnostic mode, a warning appears. <pre>filerX> priv set diag Warning: These diagnostic commands are for use by Network Appliance personnel only.</pre>

Step	Action
6	(Storage system) Enter the following command: <pre data-bbox="494 282 895 309">igroup set <igroup> vsa <value></pre> <igroup> is the name of the igroup for which you want to change the setting. <value> is either <i>on</i> or <i>off</i>. For HP-UX systems, you must always set this value to on. Result: The igroup you specify in this command. Note _____ Although this command lets you change the <i>vsa</i> value of a specific igroup and all igroups you create in the future, it does not change the value of any other existing igroups. You must perform this command on each existing igroup to change its attribute value. Example: Here is an example showing the output you might see when you change the igroup <i>hostA</i> on <i>filerX</i> to on: <pre data-bbox="494 847 1131 968">filerX> igroup set -f netapp vsa on Wed Sep 3 21:13:59 GMT [filerX: scsitarget.notice:notice]: VSA setting changed on initiator group "hostA" to "on"</pre>

Step	Action
7	(Storage system) Confirm that the <code>vsa</code> value for that igroup is now correct by using the <code>igroup show -v</code> command. ◆ If the attribute is not set correctly for an HP-UX system, the line “Volume Set Addressing: No” appears in the output. ◆ If the attribute is correctly set, there is no reference to volume set addressing or <code>vsa</code> in the output. Example: This example checks the settings for the igroup <code>hostA</code> and corrects the value for <code>vsa</code>. <pre>filerX> igroup show -v hostA hostA (FCP): OS Type: hpux Volume Set Addressing: No Member: 50:06:0b:00:00:10:a7:00 Member: 50:06:0b:00:00:10:a6:06 filerX> priv set diag Warning: These diagnostic commands are for use by Network Appliance personnel only. filerX> igroup set hostA vsa on Wed Sep 3 21:13:59 GMT [filerX: scsitarget.notice:notice]: VSA setting changed on initiator group "hostA" to "on" filerX> igroup show -v hostA hostA (FCP) (ostype: hpux): OS Type: hpux Member: 50:06:0b:00:00:10:a7:00 Member: 50:06:0b:00:00:10:a6:06 filerX> priv set</pre>

Step	Action
8	(Host) Return to the host and execute the <code>ioscan -fnC disk</code> command and the <code>ioinit -i</code> command to have the host discover the disks and create device special files for them. <pre># ioscan -fnC disk # ioinit -i</pre> Note_____ If the <code>ioinit -i</code> command does not create the devices, use the <code>insf -e</code> command. Result: The old disk devices then show up as NO_HW in the <code>ioscan</code> output.
9	(Host) Either manually remove the old device paths or reboot the host. Doing this removes the old disk devices and causes the host to discover the new ones. Make sure that the only device special files that are removed are those that disappeared as a result of the <code>igroup</code> attribute change.
10	(Host) Run the <code>vgscan -v</code> command and the <code>sanlun lun show -p</code> command to view the new locations: <pre># vgscan -v # sanlun lun show -p</pre> Correlate the new <code>sanlun</code> output with the <code>sanlun</code> information from Step 1 to determine which devices go with which volume group.

Step	Action
11	(Host) Re-add the volume group entries to <code>/etc/lvmtab</code> and the associated device files to the system. For each volume group, perform the following tasks: 1. Create a new directory with the <code>mkdir</code> command: <pre># mkdir /dev/<vg_name></pre> 2. Create a group file in the preceding directory with the <code>mknod</code> command: <pre># mknod /dev/<vg_name>/group c 64 <vg minor dev num></pre> 3. Issue the <code>vgimport</code> command with a device that represents a primary path (as indicated by the <code>sanlun lun show -p</code> command) for the volume group: <pre># vgimport /dev/<vg_name> <dev_name_primary_path></pre>
12	(Host) For each volume group you exported before changing the <code>vsa</code> attribute, activate the newly imported volume group: <pre># vgchange -a y /dev/<vg_name></pre>
13	(Host) Run the <code>ntap_config_paths</code> utility to configure multipathing for the newly imported volume groups: <pre># ntap_config_paths</pre>
14	(Host) Verify that the paths are now set up correctly: <pre># sanlun lun show all -p</pre>
15	Repeat these steps for each storage system in the cluster.

Commands you can use to troubleshoot problems

Commands you can use to check for problems

There are several commands you can use to check your system and look for problems. The following commands are especially useful in tracking down problems:

- ◆ The HP-UX command `ioscan`
- ◆ The NetApp `sanlun` utility

Using the HP-UX `ioscan` command

The HP-UX `ioscan` command displays information about the state of the LUNs and whether the host recognizes them.

Example of `ioscan` output

The following example shows the type of output from the `ioscan` command that you would see after a takeover occurs. From this example, you can determine the following information:

- ◆ There are no device entries for the NetApp LUNs yet.
- ◆ The state `CLAIMED` on two paths to the LUN means that these paths are visible to the host and ready to use.
- ◆ The state `NO_HW` on two paths to the LUN means that these paths are not visible to the host, so they cannot be used. These two paths would have been the primary paths, which were lost during the takeover.

# ioscan -fnC disk								
Class	I	H/W Path	Driver	S/W State	H/W Type	Description		
=====								
disk	0	0/0/1/1.15.0	sdisk	CLAIMED	DEVICE	HP 18.2GMAN3184MC		
			/dev/dsk/clt15d0		/dev/rdisk/clt15d0			
disk	1	0/4/0/0.1.5.0.26.0.5	sdisk	CLAIMED	DEVICE	NETAPP	LUN	
			/dev/dsk/c31t0d5		/dev/rdisk/c31t0d5			
disk	3	0/4/0/0.1.6.0.26.0.5	sdisk	NO_HW	DEVICE	NETAPP	LUN	
			/dev/dsk/c34t0d5		/dev/rdisk/c34t0d5			
disk	2	0/6/2/0.1.13.0.26.0.5	sdisk	CLAIMED	DEVICE	NETAPP	LUN	
			/dev/dsk/c30t0d5		/dev/rdisk/c30t0d5			
disk	4	0/6/2/0.1.14.0.26.0.5	sdisk	NO_HW	DEVICE	NETAPP	LUN	
			/dev/dsk/c35t0d5		/dev/rdisk/c35t0d5			

Using NetApp's sanlun utility

You can use NetApp's sanlun utility to perform the following tasks:

- ◆ Check the status of the
 - ❖ HBA
 - ❖ LUN
 - ❖ PVlinks path
- ◆ Determine the /dev/dsk/<filename> association with the storage system LUN
- ◆ Discover whether a path is up or down

With the `-p` option, you can also use the sanlun utility to

- ◆ Determine whether the path to the filer is primary or secondary.
- ◆ Report missing paths from the volume group. (You can use the `ntap_config_paths` utility to easily correct any missing paths.)
- ◆ Report the PVlinks path ordering.
- ◆ Get information about a takeover.

Using sanlun to check device special files

Using the `-pvv` option with `sanlun` increases the amount of information you get about a system. The following example illustrates how you can use `sanlun` with this option to determine which device special files are not responding to which commands:

```
# sanlun lun show -pvv

[open_device_filename] /dev/rdisk/c31t0d5
OK [/dev/rdisk/c31t0d5]
[open_device_filename] /dev/rdisk/c34t0d5
OK [/dev/rdisk/c34t0d5]
[open_device_filename] /dev/rdisk/c30t0d5
OK [/dev/rdisk/c30t0d5]
[open_device_filename] /dev/rdisk/c35t0d5
OK [/dev/rdisk/c35t0d5]

filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
      2g (2147483648))          lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy: A/P
Multipath-provider: HP PVlinks
-----
host      filer                                primary partner PVlinks
path      path      /dev/dsk      host filer      filer      path failover
state     type      filename      HBA  port      port      priority
-----
up        primary    /dev/dsk/c35t0d5 td1  5a                0
up        primary    /dev/dsk/c34t0d5 td0  4a                1
up        secondary   /dev/dsk/c30t0d5 td1                5b                2
up        secondary   /dev/dsk/c31t0d5 td0                4b                3
```

Using sanlun to check multipathing

The order of the paths you use when you set up the LUNs can affect performance. In general, secondary paths are slower than primary paths.

To get the best performance, you must set up your paths with the primary paths first, followed by the secondary paths. You can use `sanlun` with the `-p` option to check your paths. If the paths are not in the correct order, `sanlun` puts the word (wrong) next to them.

Example: The following example shows where the secondary paths were entered before the primary path. Also, in this example, another primary path is visible, but is not in the volume group.

```
# sanlun lun show -p

filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
      2g (2147483648)          lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy: A/P
Multipath-provider: HP PVlinks
-----
```

host	filer		primary	partner	PVlinks	
path	path	/dev/dsk	host	filer	filer	
state	type	filename	HBA	port	port	
					path failover	
					priority	
up	secondary	/dev/dsk/c30t0d5	td1		5b	0 (Wrong)
up	secondary	/dev/dsk/c31t0d5	td0		4b	1 (Wrong)
up	primary	/dev/dsk/c35t0d5	td1	5a		2
up	primary	/dev/dsk/c34t0d5	td0	4a		not in VG

```
#
```

You can repair the ordering of the paths by running the `ntap_config_paths` utility.

Example of sanlun output after storage system takeover

The following example shows the type of output the `sanlun` utility displays after a takeover occurs. From this example, you can determine

- ◆ The two primary paths to the LUN are down.
- ◆ The two secondary paths are taking over. The takeover paths start with the active path that has the lowest priority.

```
# sanlun lun show -p

filer1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/vg_ntap01
  2g (2147483648)    lun state: GOOD
Filer_CF_State: Partner Takeover Multipath_Policy: A/P
Multipath-provider: HP PV Links
-----
host      filer                                primary partner PVlinks
path      path                               /dev/dsk host filer  filer  path failover
state     type                               filename HBA  port   port   priority
-----
DOWN      /dev/dsk/c31t0d5                                0
DOWN      /dev/dsk/c34t0d5                                1
up        secondary /dev/dsk/c30t0d5 fcd1          0d     2
up        secondary /dev/dsk/c31t0d5 fcd0          0d     3
```

NetApp's diagnostic utilities

NetApp's diagnostic utilities

NetApp provides several scripts 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 create a gzipped .tar file and place the information in it. You can then send this information to NetApp.

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

- ◆ filer_info. See “[Collecting filer information](#)” on page 109 for more information.
- ◆ hpux_info. See “[Using the HP-UX data-collection tool](#)” on page 111.
- ◆ <switch>_info, where <switch> is the name of the switch you are using (for example, brocade_info). See “[Collecting switch information](#)” on page 111.

General steps for executing the programs

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

Step	Action
1	Go to the directory /opt/NetApp/santools/bin: <code>cd /opt/NetApp/santools/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 gzipped .tar file.

Collecting filer information

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

Note

This utility runs in command-line mode only.

Step	Action
1	Go to the directory /opt/NetApp/santools/bin: <pre>cd /opt/NetApp/santools/bin</pre>
2	At the host command prompt, enter the following command: <pre>filer_info [-d <path_and_file_name>] [-l <user:password>] -n <filer_name></pre> The -d option enables you to specify a different path in which to place the resulting output. The default path is /tmp/NetApp/ntap_filer_<filer_name>. The -l <user> option allows you to enter a user name and password on the storage system. The default login is root. The target of the command is specified by “-n <filer_name>”, where <i>filer_name</i> is the name of IP address of the filer to be interrogated. Result: The output of various status commands are placed in the output directory named ntap_filer_<filer_name>, where <i>filer_name</i> is replaced by the filer name or IP address. The contents of this directory are placed in a tarred and gzipped file named ntap_filer_<filer_name>.tar.gz. Example: The following command line checks filer1, logging on to it as the user root. <pre>host1> filer_info -d /opt/NetApp/santools -l root -n filer1</pre> The information from filer1 is placed in a gzipped file with the following directory path: <pre>/opt/NetApp/santools/ntap_filer_filer1.tar.gz</pre>
3	Send the resulting gzipped file to customer support for analysis.

**Using the HP-UX
data-collection tool**

If you are having a problem, customer support might ask you to run `hpux_info` to collect configuration information on your HP-UX system. You should run this utility only when requested to do so by NetApp.

To execute `hpux_info`, complete the following steps.

Step	Action
1	Change to the <code>/opt/NetApp/santools/bin</code> directory: <code>cd /opt/NetApp/santools/bin</code>
2	At the command prompt, enter the following command: <code>hpux_info [-d <directory>] [-n <file_name>]</code> <code>-d <directory></code> specifies the location for the gzipped file that <code>hpux_info</code> creates. The default location is <code>/tmp/NetApp</code> . <code>-n <name></code> specifies the name for the output file and directory names. The default path name is <code>/tmp/NetApp/ntap_hpux_info.tar.gz</code> . Result: When you run <code>hpux_info</code> , the utility gathers information about your HP-UX configuration and places it in a gzipped file. It also creates a directory and populates it with files containing output from numerous commands and copies of files, such as <code>/var/adm/syslog/syslog.log</code> file. Example: Depending on the options you specify, you see output similar to the following when you run this utility: <pre># hpux_info HP-UX system info is in directory /tmp/NetApp/ntap_hpux_info Compressed file is /tmp/NetApp/ntap_hpux_info.tar.gz Please send this file to customer support for analysis</pre>
3	Send the resulting gzipped file to customer support for analysis.

**Collecting switch
information**

If you are having a problem, customer support might ask you to run the switch utility NetApp supplies when you install the host Host Utilities software. NetApp provides a utility for each supported switch for this Host Utilities kit. The utility name is `<switch>_info`, where `switch` is the name of the switch (for example, `brocade_info` or `mcdata_info`). This utility collects switch configuration information. You should run this utility only when requested to do so by NetApp.

To execute the utility, complete the following steps.

Note

This utility runs in command-line mode only. Double-clicking the utility name to run it does not work.

Step	Action
1	Go to the directory /opt/NetApp/santools/bin: <code>cd /opt/NetApp/santools/bin</code>

Step	Action
2	At the host command prompt, enter the following command: <pre><switch_brand>_info [-d <path_and_file_name>] [-l <user:password>] -n <switchname></pre> <switch_brand> is the name of the switch you are using; for example, brocade or mcddata. The -d <path_and_file_name> option enables you to specify a different path in which to place the resulting output. The default path is //tmp/NetApp/<switch_brand>_switch_<switch_name>. The -l <user:password> 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. The target of the command is specified by -n <switchname>, where <i>switchname</i> is the name or IP address of the switch to be interrogated. Result: The output of various status commands on the switch are placed in a directory named <switch_brand>_switch_<switchname>, 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 tarred and gzipped file named <switch_brand>_switch_<switchname>.tar.gz. Example: This example uses the brocade_info file. The switch brcd1 is connected with the user "john", who has the password "smith". <pre>host1> brocade_info -d /opt/NetApp/santools -l john:smith -n brcd1</pre> The information from filer1 is placed in a tarred, gzipped file that has the pathname /opt/NetApp/santools/brocade_switch_brcd1.tar.gz.
3	Send the resulting gzipped file to customer support for analysis.

General troubleshooting

Basic operations take a long time

Problem: There are no errors and everything appears to be correctly connected, but basic operations are taking exceptionally long times; in some cases, hours instead of minutes.

Explanation: If there are no other problems, this could mean that one of the connectors is not firmly attached. This can be the case even when it has a link and is reporting information such as the correct speed.

What to do: Disconnect the connector and reconnect it.

Appendix overview

If you are using Veritas Storage Foundation for multipathing, you must install and configure the Symantec Array Support Library (ASL) for NetApp storage systems. This appendix describes how to install and configure the ASL and how this component affects multipathing and failover.

Note

To determine which version of the Veritas Storage Foundation and ASL to use, check the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Topics in this appendix

Topics in this appendix include the following:

- ◆ “[Overview of the Array Support Library](#)” on page 116
- ◆ “[Installing the ASL](#)” on page 117
- ◆ “[Removing the ASL](#)” on page 119
- ◆ “[Obtaining multipathing information](#)” on page 120
- ◆ “[Understanding ASL error message types](#)” on page 125

Overview of the Array Support Library

What the Array Support Library is

The ASL is a NetApp-qualified library that provides information about storage array attributes and multipathing configuration to the Device Discovery Layer (DDL) and Dynamic Multipathing (DMP) components of 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. The ASL in turn “claims” a device based on vendor and product identifiers. The claim associates the storage array model and product identifiers with the device.

The ASL supports VxVM’s enclosure-based naming feature. This means the name of the disk is based on the logical name of its enclosure, or disk array. The ASL provides specific vendor and model information to DMP and VxVM instead of referring to them as JBOD devices or raw devices.

Note

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.

You can install or remove the ASL without rebooting the host.

Note

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

What the ASL provides

The ASL identifies which LUNs are NetApp storage so that these LUNs are not identified as JBOD. By using the ASL, you obtain the following information about the LUNs:

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

Installing the ASL

ASL installation overview

Installing the ASL involves the following tasks:

- ◆ Determining the ASL version.
- ◆ Obtaining the ASL.
- ◆ Verifying that your configuration meets system requirements. See the *Compatibility and Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for current information about the system requirements.

Obtaining the ASL software

The ASL is not distributed with the Host Utilities. You obtain the ASL from the Symantec Web site. Download this software before you start the installation.

For a direct link to the ASL on the Symantec Web site, see the FCP Host Compatibility Matrix at

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

Note

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

Installing the ASL software

To install a fresh version of the ASL from Symantec, complete the following steps.

Step	Action
1	Log in to the VxVM system as the root user.
2	If you already have NetApp storage configured as JBOD in your VxVM configuration, remove the JBOD support for NetApp by entering: vxddladm rmjbof vid=NETAPP
3	Install the ASL according to the directions provided in the ASL TechNote on the Symantec Web site. You should have your LUNs set up before you install the ASL. For information on setting up LUNs, see “ Creating, Configuring, and Managing LUNs ” on page 21.

Step	Action																																																												
4	If your host is connected to NetApp storage and you can verify the installation by entering the following command and then locating the NetApp Enclosure Type in the output of the <code>vxddmpadm listenclosure all</code> command. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM. <pre># vxddmpadm listenclosure all</pre> <table><thead><tr><th>ENCLR_NAME</th><th>ENCLR_TYPE</th><th>ENCLR_SNO</th><th>STATUS</th><th>ARRAY_TYPE</th></tr></thead><tbody><tr><td colspan="5">=====</td></tr><tr><td>Disk</td><td>Disk</td><td>DISKS</td><td>CONNECTED</td><td>Disk</td></tr><tr><td>FAS9800</td><td>FAS980</td><td>110296811</td><td>CONNECTED</td><td>A/PG-C</td></tr><tr><td>FAS9801</td><td>FAS980</td><td>110296822</td><td>CONNECTED</td><td>A/PG-C</td></tr></tbody></table>	ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE	=====					Disk	Disk	DISKS	CONNECTED	Disk	FAS9800	FAS980	110296811	CONNECTED	A/PG-C	FAS9801	FAS980	110296822	CONNECTED	A/PG-C																																			
ENCLR_NAME	ENCLR_TYPE	ENCLR_SNO	STATUS	ARRAY_TYPE																																																									
=====																																																													
Disk	Disk	DISKS	CONNECTED	Disk																																																									
FAS9800	FAS980	110296811	CONNECTED	A/PG-C																																																									
FAS9801	FAS980	110296822	CONNECTED	A/PG-C																																																									
5	If your host is not yet connected to storage, enter the following command: <pre>vxddladm listsupport all</pre> Example: The following is a sample of the output you see when you enter the <code>vxddladm</code> command. <pre># vxddladm listsupport all</pre> <table><thead><tr><th>LIBNAME</th><th>VID</th><th>PID</th></tr></thead><tbody><tr><td colspan="3">=====</td></tr><tr><td>libvxauid.sl</td><td>HP</td><td>C3586A, C5447A, A5257A</td></tr><tr><td>libvxcscovrts.sl</td><td>CSCOVRTS</td><td>MDS9</td></tr><tr><td>libvxdgc.sl</td><td>DGC</td><td>All</td></tr><tr><td>libvxeccs.sl</td><td>ECCS</td><td>All</td></tr><tr><td>libvxemc.sl</td><td>EMC</td><td>SYMMETRIX</td></tr><tr><td>libvxfc60.sl</td><td>HP</td><td>A5277A</td></tr><tr><td>libvxfujitsu.sl</td><td>FUJITSU</td><td>GR710, GR720, GR730</td></tr><tr><td>libvxhds.sl</td><td>HITACHI</td><td>All</td></tr><tr><td>libvxhdsalua.sl</td><td>HITACHI</td><td>DF600F</td></tr><tr><td>libvxhitachi.sl</td><td>HITACHI</td><td>DF350, DF400, DF400F</td></tr><tr><td>libvxhpalua.sl</td><td>HP</td><td>HSV210</td></tr><tr><td>libvxnec.sl</td><td>NEC</td><td>DS1200, DS1200F, DS3000SL</td></tr><tr><td>libvxshark.sl</td><td>IBM</td><td>2105</td></tr><tr><td>libvxstorcomp.sl</td><td>StorComp</td><td>OmniForce</td></tr><tr><td>libvxva.sl</td><td>HP</td><td>A6188A, A6189A</td></tr><tr><td>libvxveritas.sl</td><td>VERITAS</td><td>All</td></tr><tr><td>libvxxp256.sl</td><td>HP</td><td>All</td></tr><tr><td>libvxnetapp.sl</td><td>NETAPP</td><td>All</td></tr></tbody></table>	LIBNAME	VID	PID	=====			libvxauid.sl	HP	C3586A, C5447A, A5257A	libvxcscovrts.sl	CSCOVRTS	MDS9	libvxdgc.sl	DGC	All	libvxeccs.sl	ECCS	All	libvxemc.sl	EMC	SYMMETRIX	libvxfc60.sl	HP	A5277A	libvxfujitsu.sl	FUJITSU	GR710, GR720, GR730	libvxhds.sl	HITACHI	All	libvxhdsalua.sl	HITACHI	DF600F	libvxhitachi.sl	HITACHI	DF350, DF400, DF400F	libvxhpalua.sl	HP	HSV210	libvxnec.sl	NEC	DS1200, DS1200F, DS3000SL	libvxshark.sl	IBM	2105	libvxstorcomp.sl	StorComp	OmniForce	libvxva.sl	HP	A6188A, A6189A	libvxveritas.sl	VERITAS	All	libvxxp256.sl	HP	All	libvxnetapp.sl	NETAPP	All
LIBNAME	VID	PID																																																											
=====																																																													
libvxauid.sl	HP	C3586A, C5447A, A5257A																																																											
libvxcscovrts.sl	CSCOVRTS	MDS9																																																											
libvxdgc.sl	DGC	All																																																											
libvxeccs.sl	ECCS	All																																																											
libvxemc.sl	EMC	SYMMETRIX																																																											
libvxfc60.sl	HP	A5277A																																																											
libvxfujitsu.sl	FUJITSU	GR710, GR720, GR730																																																											
libvxhds.sl	HITACHI	All																																																											
libvxhdsalua.sl	HITACHI	DF600F																																																											
libvxhitachi.sl	HITACHI	DF350, DF400, DF400F																																																											
libvxhpalua.sl	HP	HSV210																																																											
libvxnec.sl	NEC	DS1200, DS1200F, DS3000SL																																																											
libvxshark.sl	IBM	2105																																																											
libvxstorcomp.sl	StorComp	OmniForce																																																											
libvxva.sl	HP	A6188A, A6189A																																																											
libvxveritas.sl	VERITAS	All																																																											
libvxxp256.sl	HP	All																																																											
libvxnetapp.sl	NETAPP	All																																																											

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.

Only one version of an ASL that supports Data ONTAP storage systems can be installed on the host at any given time. You cannot have multiple versions of the NetApp ASL installed simultaneously.

Removing the ASL

To remove the ASL, complete the following steps.

Step	Action
1	Log into the host as the root user.
2	Stop all I/O to NetApp LUNs. Note_____ In a Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL. _____
3	Enter the following command to remove the software: <code>swremove NTAPas1</code>

Obtaining multipathing information

What an ASL array type is

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

- ◆ **Active/Active (A/A)**—There are multiple active paths to storage and simultaneous I/O is supported on all paths. If a path fails, I/O is distributed across the remaining paths.
- ◆ **Active/Passive Group Concurrent (A/PG-C)** - An A/PG-C array is a variant of the A/PG array type that supports concurrent I/O and load balancing by having multiple primary paths to LUNs. Failover to the secondary (passive) path occurs only if all the active primary paths fail.

For additional information, see the following documents:

- ◆ *Veritas Array Support Library Developer's Guide* - for details about load balancing.
- ◆ *Veritas Volume Manager Administrator's Guide* - for details about system management

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

In clustered NetApp storage configurations, the A/PG-C array type corresponds to all storage system cfmode settings that are supported in configurations with HP-UX hosts. The standby cfmode setting is not supported in configurations with HP-UX hosts.

Note

The ASL also supports direct attached, nonclustered configurations, including NearStore models. These configurations have no cfmode settings. ASL reports these configurations as Active/Active (A/A) array types. For more details about array types, see the *Veritas Array Support Library Developer's Guide*.

How sanlun displays the array type

The FCP HP-UX Host Utilities SAN Toolkit includes the `sanlun` utility in the `/opt/NetApp/santools/bin` directory. You use the `sanlun` utility to display information about paths to LUNs on the storage system.

The FCP HP-UX Host Utilities supports both VxVM and the HP-UX LVM.

- ◆ When the ASL is installed and the LUN is controlled by VxVM, the output of the `sanlun` command displays the `Multipath_Policy` as A/PG-C for clustered storage and A/A for non-clustered (single filer) configurations. Veritas DMP is the Multipathing-provider.

- ◆ If the LUN is controlled by the HP-UX LVM, the output displays the Multipath_Policy as A/P and HP PVlinks is the Multipathing-provider.
- ◆ If the LUN is not controlled by a volume manager, then the Multipathing_Policy is none and the Multipathing-provider is none.

The following example shows partial output of the `san lun` command when a LUN is controlled by VxVM and the host is attached to a NetApp cluster.

```
filerX:/vol/vol1/vxvm_data1 (LUN 1) DMP-NODE: (F880_1)
      25m (26214400)          lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipathing_Policy: A/PG
Multipath-provider: Veritas DMP
```

Using VxVM to display available paths

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

Step	Action
1	Enter the following command to view all devices: <pre>vxdisk list</pre> Result: The VxVM management interface displays the vxdisk device, type, disk, group and status. It also shows which devices are managed by the HP-UX Logical Volume Manager (LVM) and which disks are managed by VxVM. Example: The following is an example of the type of output you see with the vxdisk list command. <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:LVM - - LVM Disk_1 auto:LVM - - LVM FAS9800_0 auto:cdsdisk DG-2-disk19 DG-2 online FAS9800_1 auto:cdsdisk DG-2-disk18 DG-2 online FAS9800_2 auto:cdsdisk DG-2-disk20 DG-2 online FAS9800_3 auto:cdsdisk DG-2-disk17 DG-2 online FAS9800_4 auto:cdsdisk DG-2-disk16 DG-2 online FAS9800_5 auto:cdsdisk DG-2-disk15 DG-2 online FAS9800_6 auto:cdsdisk DG-2-disk12 DG-2 online FAS9800_7 auto:cdsdisk DG-2-disk13 DG-2 online FAS9800_8 auto:cdsdisk DG-2-disk14 DG-2 online FAS9800_9 auto:cdsdisk DG-2-disk11 DG-2 online FAS9801_0 auto:cdsdisk DG-2-disk10 DG-2 online FAS9801_1 auto:cdsdisk DG-2-disk6 DG-2 online FAS9801_2 auto:cdsdisk DG-2-disk7 DG-2 online FAS9801_3 auto:cdsdisk DG-2-disk5 DG-2 online FAS9801_4 auto:cdsdisk DG-2-disk8 DG-2 online FAS9801_5 auto:cdsdisk DG-2-disk3 DG-2 online FAS9801_6 auto:cdsdisk DG-2-disk1 DG-2 online FAS9801_7 auto:cdsdisk DG-2-disk9 DG-2 online FAS9801_8 auto:cdsdisk DG-2-disk4 DG-2 online FAS9801_9 auto:cdsdisk DG-2-disk2 DG-2 online</pre>

Step	Action																																																																																				
2	On the host console, enter the following command to display path information for the device you want: <pre>vxddmpadm getsubpaths dmpnodename=<device></pre> where <device> is the name listed under the output of the <code>vxdisk list</code> command. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths dmpnodename=FAS9800_0</pre> <table><tr><th>NAME</th><th>STATE [A]</th><th>PATH-TYPE [M]</th><th>CTLR-NAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td colspan="7">=====</td></tr><tr><td>c6t1d0</td><td>ENABLED (A)</td><td>PRIMARY</td><td>c6</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c10t1d0</td><td>ENABLED</td><td>SECONDARY</td><td>c10</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c16t1d0</td><td>ENABLED (A)</td><td>PRIMARY</td><td>c16</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c20t1d0</td><td>ENABLED</td><td>SECONDARY</td><td>c20</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr></table>	NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====							c6t1d0	ENABLED (A)	PRIMARY	c6	FAS980	FAS9800	-	c10t1d0	ENABLED	SECONDARY	c10	FAS980	FAS9800	-	c16t1d0	ENABLED (A)	PRIMARY	c16	FAS980	FAS9800	-	c20t1d0	ENABLED	SECONDARY	c20	FAS980	FAS9800	-																																										
NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																															
=====																																																																																					
c6t1d0	ENABLED (A)	PRIMARY	c6	FAS980	FAS9800	-																																																																															
c10t1d0	ENABLED	SECONDARY	c10	FAS980	FAS9800	-																																																																															
c16t1d0	ENABLED (A)	PRIMARY	c16	FAS980	FAS9800	-																																																																															
c20t1d0	ENABLED	SECONDARY	c20	FAS980	FAS9800	-																																																																															
3	Enter the following command to obtain path information for a host HBA: <pre>vxddmpadm getsubpaths ctlr=<controller_name></pre> controller_name is the controller displayed under “CTRL - NAME” in the output of the <code>vxddmpadm getsubpaths dmpnodename</code> command you entered in Step 2. Result: The output displays information about the paths to the filer storage (whether the path is a primary or secondary path). The output also lists the filer that the device is mapped to. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths ctlr=c6</pre> <table><tr><th>NAME</th><th>STATE [A]</th><th>PATH-TYPE [M]</th><th>DMPNODENAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td colspan="7">=====</td></tr><tr><td>c6t0d0</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_9</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d1</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_6</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d2</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_7</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d3</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_8</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d4</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_5</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d5</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_4</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d6</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_3</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t0d7</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_1</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t1d0</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_0</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr><tr><td>c6t1d1</td><td>ENABLED (A)</td><td>PRIMARY</td><td>FAS9800_2</td><td>FAS980</td><td>FAS9800</td><td>-</td></tr></table>	NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====							c6t0d0	ENABLED (A)	PRIMARY	FAS9800_9	FAS980	FAS9800	-	c6t0d1	ENABLED (A)	PRIMARY	FAS9800_6	FAS980	FAS9800	-	c6t0d2	ENABLED (A)	PRIMARY	FAS9800_7	FAS980	FAS9800	-	c6t0d3	ENABLED (A)	PRIMARY	FAS9800_8	FAS980	FAS9800	-	c6t0d4	ENABLED (A)	PRIMARY	FAS9800_5	FAS980	FAS9800	-	c6t0d5	ENABLED (A)	PRIMARY	FAS9800_4	FAS980	FAS9800	-	c6t0d6	ENABLED (A)	PRIMARY	FAS9800_3	FAS980	FAS9800	-	c6t0d7	ENABLED (A)	PRIMARY	FAS9800_1	FAS980	FAS9800	-	c6t1d0	ENABLED (A)	PRIMARY	FAS9800_0	FAS980	FAS9800	-	c6t1d1	ENABLED (A)	PRIMARY	FAS9800_2	FAS980	FAS9800	-
NAME	STATE [A]	PATH-TYPE [M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																															
=====																																																																																					
c6t0d0	ENABLED (A)	PRIMARY	FAS9800_9	FAS980	FAS9800	-																																																																															
c6t0d1	ENABLED (A)	PRIMARY	FAS9800_6	FAS980	FAS9800	-																																																																															
c6t0d2	ENABLED (A)	PRIMARY	FAS9800_7	FAS980	FAS9800	-																																																																															
c6t0d3	ENABLED (A)	PRIMARY	FAS9800_8	FAS980	FAS9800	-																																																																															
c6t0d4	ENABLED (A)	PRIMARY	FAS9800_5	FAS980	FAS9800	-																																																																															
c6t0d5	ENABLED (A)	PRIMARY	FAS9800_4	FAS980	FAS9800	-																																																																															
c6t0d6	ENABLED (A)	PRIMARY	FAS9800_3	FAS980	FAS9800	-																																																																															
c6t0d7	ENABLED (A)	PRIMARY	FAS9800_1	FAS980	FAS9800	-																																																																															
c6t1d0	ENABLED (A)	PRIMARY	FAS9800_0	FAS980	FAS9800	-																																																																															
c6t1d1	ENABLED (A)	PRIMARY	FAS9800_2	FAS980	FAS9800	-																																																																															

Using sanlun to obtain multipathing information

To obtain multipathing information by using the `san lun` command, complete the following step.

Step	Action																																																																																																												
1	On the host command line, enter the following command: <pre>sanlun lun show -p all</pre> Result: san lun displays path information for each LUN. For LUNs managed by VxVM, the output displays the Veritas multipathing policy (A/A or A/PG-C array types) and the path I/O policy. For LUNs managed by LVM, the output displays the PVlinks path order. Example: The following partial output shows paths to two LUNs managed by VxVM. Note that the output shows the Veritas path I/O policy as Not Applicable. <pre># sanlun lun show -p all filerX:/vol/vol1/vxvm_data1 (LUN 1) DMP-NODE: (F880_1) 25m (26214400) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG Multipath-provider: Veritas DMP</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>Veritas</th></tr><tr><th>path</th><th>path</th><th>/dev/dsk</th><th>host</th><th>filer</th><th>filer</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>path I/O</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>policy</th></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c17t0d1</td><td>td1</td><td>5a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c16t0d1</td><td>td1</td><td></td><td>5b</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c19t0d1</td><td>td0</td><td>4a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c18t0d1</td><td>td0</td><td></td><td>4b</td></tr></table> <pre>filerX:/vol/test/lun-0 (LUN 0) DMP-NODE: (FAS9801_6) 10g (10737418240) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG Multipath-provider: Veritas DMP</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>Veritas</th></tr><tr><th>path</th><th>path</th><th>/dev/dsk</th><th>host</th><th>filer</th><th>filer</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>path I/O</th></tr><tr><th></th><th></th><th></th><th></th><th></th><th>policy</th></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c22t0d0</td><td>td1</td><td>8a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c18t0d0</td><td>td1</td><td></td><td>8b</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c8t0d0</td><td>td0</td><td>9a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c4t0d0</td><td>td0</td><td></td><td>9b</td></tr></table>	host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	primary	/dev/dsk/c17t0d1	td1	5a		up	secondary	/dev/dsk/c16t0d1	td1		5b	up	primary	/dev/dsk/c19t0d1	td0	4a		up	secondary	/dev/dsk/c18t0d1	td0		4b	host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	primary	/dev/dsk/c22t0d0	td1	8a		up	secondary	/dev/dsk/c18t0d0	td1		8b	up	primary	/dev/dsk/c8t0d0	td0	9a		up	secondary	/dev/dsk/c4t0d0	td0		9b
host	filer		primary	partner	Veritas																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	primary	/dev/dsk/c17t0d1	td1	5a																																																																																																									
up	secondary	/dev/dsk/c16t0d1	td1		5b																																																																																																								
up	primary	/dev/dsk/c19t0d1	td0	4a																																																																																																									
up	secondary	/dev/dsk/c18t0d1	td0		4b																																																																																																								
host	filer		primary	partner	Veritas																																																																																																								
path	path	/dev/dsk	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	primary	/dev/dsk/c22t0d0	td1	8a																																																																																																									
up	secondary	/dev/dsk/c18t0d0	td1		8b																																																																																																								
up	primary	/dev/dsk/c8t0d0	td0	9a																																																																																																									
up	secondary	/dev/dsk/c4t0d0	td0		9b																																																																																																								

Understanding ASL error message types

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

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

Index

Symbols

/etc/MANPATH 15

/etc/PATH 15

/opt/NetApp/santools/man 15

A

A/A

Active/Active paths 120

A/PG-C

Active/Passive Group Concurrent paths 120

active paths

Active/Active (A/A) 120

Active/Passive Group Concurrent paths

A/PG-c 120

Array Support Library

see ASL 116

ASL

components 116

definition 116

installing 117

multipathing reports 120

removing 119

Attach Kits

previous kits 2

C

cfmode

check setting 23

checklist

filer 88

host 89

troubleshooting 88

checklists

installation tasks 5

commands

gunzip 13

igroup create 25

igroup show 26

ioinit (HP-UX) 37

ioscan (HP-UX) 36

lun comment 29

lun setup 26, 27

lun show 34

lvcreate (HP-UX) 42

mount (HP-UX) 44

nwfs (HP-UX) 44

sanlun fcp show adapter 57

sanlun lun show 38, 45, 53

swremove 17

vgcreate (HP-UX) 40

vgdisplay (HP-UX) 42

D

driver

logs off target port 88

E

enter

defined vii

examples

checking cfmode 23

F

FCP HP-UX Host Utilities

contents 2

defined 2

obtaining 13

SAN Toolkit 3, 13

software 13

file system

create 44

mount 44

filer

adding LUNs 22

defined vii

entering commands viii

setup checklist 88

filer_info

diagnostic script 3

G

- gunzip command
 - uncompressing NetApp file 13

H

HBA

- displaying information 57

host

- adding LUNs 22
- create file system 44
- discovering LUNs 35
- displaying HBA information 57
- displaying LUN information 53
- setup checklist 89

HP-UX

- connectivity issues 88
- FCP driver behavior 88
- LVM 35

I

I/O

- numerous retries 88

igroup

- create command 25
- creating 24
- definition 24
- verify 26
- vsa attribute 26

- igroup create command 25

- igroup show command 26

- initiator group, creating 24

- initiator group

- see igroup 24

installation

- NetApp software package 13
- tasks performed 12
- update to search paths 15

- ioinit command 37

- ioscan command 36

L

- logcial volumes

- creating 42

- Logical Volume Manager 2

LUN

- defined vii

- lun commands

- lun create 64

- lun map 64

- lun comment command 29

- LUN ID 33

- defined vii

- lun setup command 26

- command line options 27

- interactive script 27

- lun show command 34

LUNs

- adding 22

- creating with command line 27

- creating with interactive script 26

- discovered by host 35

- displaying information on host 53

- lun setup command 27

- tasks for creating on filer 24

- verifying 34

- lvcreate command 42

LVM

- Logcial Volume Manager 35

- used for volume management 2

- volume group entries 39

M

- mount command 44

multipathing

- LVM 3

- VxVM 3

N

- NetApp software package

- installing 13

- ntap_config_paths utility 41

- reparing paths 92

- nwfs command 44

O

opt/NetApp/santools/bin 15

P

path order

volume groups 41

paths

displaying with sanlun lun show command 53

repairing with ntap_config_paths 92

PVLinks

uses ntap_config_paths 3

S

SAN Toolkit 3

components 3

obtaining 13

san_version command 3

sanlun utility 3

san_version

command 3

sanlun

part of SAN Toolkit 3

sanlun fcp show adapter command 57

sanlun lun show command 38, 45, 53

sanlun utility

about 52

locating 22

sanlun fcp show adapter 57

scripts

filer_info 3

ntap_config_paths 3

script used with PVLinks 3

software 13

swremove command

uninstalling software 17

system search path

updated 15

T

troubleshooting

checklists 88

type

defined vii

V

vgcreate command 40

vgdisplay command 42

volume groups

creating logical volumes 42

creating with primary path 40

fixing while offline 90

fixing while online 90

minor numbers 39

path order 40, 41

situations requiring reconfiguration 90

using ntap_conf_paths 41

vsa attribute 26

VxVM

used for volume management 2

