

FC HP-UX® Host Utilities 4.3 for Native OS and Veritas® **Installation and Setup Guide**

NetApp, 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: 215-04256_A0
September 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this file 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 material of NetApp, Inc. is subject to the following license and disclaimer:

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

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

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

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

Trademark information

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

Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, AIX, and System Storage are trademarks and/or registered trademarks of International Business Machines Corporation.

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

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

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

Table of Contents

	Preface	vii
Chapter 1	Host Utilities Overview	1
	Host Utilities Overview.	2
	Where to go for more information	5
Chapter 2	Installing Host Utilities	7
	Installation overview	9
	Installing FC HP-UX Host Utilities 4.3	11
	Removing FC HP-UX Host Utilities SAN Toolkit or Attach Kit	15
Chapter 3	Configuring SAN for FC	19
	Configuring the storage system.	20
	Configuring FC cabling and switches	22
Chapter 4	Creating, Configuring, and Managing LUNs	25
	Overview of adding LUNs to an HP-UX host	26
	Creating igroups and LUNs.	27
	Configuring the LUNs	38
	Displaying LUN and HBA information	61
	Displaying host LUN information	62
	Displaying host HBA information	67
Chapter 5	Configuring SAN Boot	69
	Overview of SAN boot procedures.	70
	Creating the boot LUN using LVM (LVM only)	73
	Migrating an OS disk using MirrorDisk/UX on PA-RISC systems (LVM only)	79
	Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems (LVM only).	87

	Testing the alternate boot path for PA-RISC systems (LVM only)	93
	Troubleshooting sanboot configurations with PA-RISC systems (LVM only)	94
	Creating the boot LUN using vxcp_lvmroot (Veritas only)	96
	Creating the boot LUN using a cold install from the HP-UX DVD (Veritas only)	97
Chapter 6	Veritas Array Support Library and Array Policy Module.	99
	Overview of the ASL and APM	100
	Installing the ASL and APM	102
	Removing the ASL and APM	105
	Obtaining multipathing information	109
	Understanding Array Support Library error message types	114
Chapter 7	Troubleshooting	115
	Troubleshooting checklists	116
	Correcting problems with Veritas	119
	Correcting volume group multipathing with ntap_config_paths	120
	Handling a takeover and giveback	122
	Recovering from a host reboot during a takeover	123
	Enabling the host to identify more than eight LUNs per target HBA	126
	Commands you can use to troubleshoot problems	132
	Diagnostic utilities	137
	General troubleshooting	142

Preface

About this guide

This guide describes how to configure an HP-UX® host in a storage area network (SAN) environment. The host accesses data on a NetApp® storage system that runs Data ONTAP® software with Fibre Channel (FC) protocol.

FC HP-UX Host Utilities 4.3 provides support for native Logical Volume Manager (LVM) on both HP-UX 11iv2 and 11iv3; it supports Veritas Storage Foundation™ on HP-UX 11iv2 only.

For information about specific HP-UX versions supported, see the one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Audience

This guide is for the following system installers and administrators:

- ◆ Who are familiar with the HP-UX operating system
- ◆ Who understand how to configure and administer a storage system
- ◆ Who understand how to use the storage-licensed SAN protocol
- ◆ Who understand how to use FC to store and retrieve data on a storage system

Terminology

This guide uses the following terms:

- ◆ NetApp storage products (storage systems, FAS appliances, and NearStore® systems) are all storage systems and are also sometimes called *storage appliances*.
- ◆ *igroup* (initiator group) is a collection of FC worldwide port names (WWPNs) of the HBA ports that have access to any LUNs mapped to the group.
- ◆ *LUN* (logical unit number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Initiator* refers to the system component that originates an I/O command over an I/O bus or network.

- ◆ *Target* refers to the storage system that receives an FC I/O command and responds to I/O request from the initiator.
- ◆ *Fabric* refers to the interconnection of two or more Fibre Channel switches in such a manner that data can be physically transmitted between any of the switches.

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>fc test adaptername</code> command, you enter the characters “ <code>fc test</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.
<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.

Attention _____
An Attention notice 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, put FC HP-UX in the subject line.

About this chapter

This chapter provides an overview of the tasks that you need to perform to install Host Utilities and set up your system. It also includes information about other documentation that can be helpful in working with your Host Utilities setup.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Host Utilities Overview”](#) on page 2
- ◆ [“ALUA support”](#) on page 4
- ◆ [“Where to go for more information”](#) on page 5

Host Utilities Overview

Important information about Host Utilities and Attach Kits

FC HP-UX Host Utilities for Native OS and Veritas® provides the tools necessary to connect HP-UX hosts to NetApp storage systems.

FC HP-UX Host Utilities supports HP-UX Logical Volume Manager (LVM) and Veritas Volume Manager (VxVM) to enable you to have different volume management and multipathing solutions.

Note

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

NetApp provides software and documentation for HP-UX Host Utilities in FC HP-UX Host Utilities 4.3. Component versions are not tied to the Host Utilities versions, so it is possible for each component to have a different version. For more information about version compatibility, see one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Overview of FC HP-UX Host Utilities : FC HP-UX Host Utilities 4.3 provides the tools and information necessary to enable an HP-UX host to connect to your storage systems.

Host Utilities consists of the following:

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

For information about the supported stack, see one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

NetApp adds components on an ongoing basis. NetApp qualifies new host bus adapters (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

FC HP-UX Host Utilities provides 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:

- ◆ **sanlun utility**
You can use this utility to display information about LUNs on the storage system that are available to this host.
- ◆ **enable_ontap_pvlinks script**
This script helps to enabling the active/passive multipathing policy for PVLINKs with Data ONTAP LUNs. This script is applicable to the HP-UX11iv3 only.
- ◆ **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 Physical Volume Links (PVLINKs), you can use `ntap_config_paths` to add missing paths to LVM's 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.
- ◆ **Additional diagnostic scripts**
The following scripts can also be used to collect troubleshooting information:
 - ❖ **brocade_info** collects information about Brocade® switches installed in the network.
 - ❖ **cisco_info** collects information about Cisco® switches installed in the network.
 - ❖ **mcdata_info** collects information about McDATA® switches installed in the network.
 - ❖ **qlogic_info** collects information about QLogic® switches installed in the network.

- ❖ `SHsupport.pm` is the support library for the `*info` scripts.
- ❖ `FTP.pm` is the support library for `* info` scripts.
- ❖ `Telnet.pm` support library for `*info` scripts.
- ◆ The manual (`man`) pages for `sanlun`, the diagnostic scripts, `enable_ontap_pvlinks` script, and the `ntap_config_paths` script.

Volume managers and multipathing

FC HP-UX Host Utilities manage the volumes using either HP-UX LVM or Veritas VxVM. In some cases you might 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 LVM's 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 SymantecTM Array Support Library (ASL) with veritas 4.1 and Symantec Array Support Library/Array Policy Module (ASL/APM) with Veritas 5.0 for NetApp with Data ONTAP Storage Foundation for HP-UX. For information, see one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

ALUA support

The FC HP-UX Host Utilities adds support for Asymmetric Logical Unit Access (ALUA). Check the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Note

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

ALUA defines a standard set of SCSI commands for discovering and managing multiple paths to LUNs on Fibre Channel and iSCSI SANs.

Where to go for more information

For host and storage system information

The following table describes the documents on the NOW® site. The documents 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 FC HP-UX Host Utilities, which is available from the NOW site.
The latest supported configurations	◆ The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml. ◆ The <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.cgi
Supported Fibre Channel SAN topologies	The NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability , or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml .
Configuring the storage system	◆ The <i>Best Practices for Reliability: New System Installation</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/index.shtml. ◆ The <i>Data ONTAP Software Setup Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
Host Operating system information or settings	http://docs.hp.com/en/native-multipathing/native_multipathing_wp_AR0709.pdf

If you need information about...	See the following NetApp product documentation...
Managing SAN storage on the storage system	◆ The Data ONTAP® <i>Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The Data ONTAP® <i>Release Notes</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The <i>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ The FilerView® <i>Online Help</i> - available from within the FilerView GUI ◆ The <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	The <i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.cgi .
Upgrading Data ONTAP	The Data ONTAP® <i>Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
Installing and configuring the HBA in your host	Your HBA vendor documentation.
Your host operating system and using its features, such as LVM and VxVM	Your operating system documentation. You can download HP-UX manuals in the PDF format from http://docs.hp.com/ .
Enabling ALUA	<i>Data ONTAP Block Access Management Guide</i> for FC at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml . See the section about enabling ALUA for a Fibre Channel igroup.

About this chapter

This chapter provides an overview of the tasks that you need to perform to install Host Utilities and set up your system. The procedures for installing the Host Utilities software package for an HP-UX host connected to a system running Data ONTAP with FC. It also includes instructions for uninstalling the software and the information about other documentation that can be helpful in working with your Host Utilities setup.

Note

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

Topics in this chapter

This chapter contains the following topics:

- ◆ “[Installation overview](#)” on page 9
- ◆ “[Installing FC HP-UX Host Utilities 4.3](#)” on page 11
- ◆ “[Removing FC HP-UX Host Utilities SAN Toolkit or Attach Kit](#)” on page 15

Pre-requisites

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

Step	Action
1	Set up your system: ◆ Host operating system (OS) and appropriate updates ◆ HBAs ◆ Drivers ◆ Storage system and Data ONTAP ◆ Switches see one of the following references: ◆ NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Step	Action
2	Verify the following for your storage system: ◆ Your storage system is licensed correctly ◆ Your storage system running the FC service ◆ ALUA is enabled, if you are using 11iv3 September 2007 release and later ◆ Storage system using the correct <code>cfmode</code> setting ◆ See one of the following reference for information about supported cluster failover modes: ❖ NetApp Interoperability Matrix at http://now.netapp.com/NOW/products/interoperability ❖ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml ◆ Your storage system configured to work with the target HBAs (if needed) ◆ Your storage system is set up to work with the HP-UX host and the initiator HBAs ◆ Your storage system is 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 are using a fabric connection, verify the following for the switch: ◆ The switch is set up correctly ◆ It is cabled according to the instructions that came with it ◆ It is zoned (if appropriate for your system)
4	Verify the following connectivity: ◆ 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 Host Utilities involves the following tasks.

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

For detailed information about installing Host Utilities

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

- ◆ [“Installing FC HP-UX Host Utilities 4.3”](#) on page 11
- ◆ [“Configuring SAN for FC”](#) on page 19
- ◆ [“Creating igroups and LUNs”](#) on page 27
- ◆ [“Configuring LUNs”](#) on page 38
- ◆ [“Displaying LUN and HBA information”](#) on page 61

In addition, the section [“Where to go for more information”](#) on page 5 provides information about additional documents and Web sites that can help you with your system setup.

Optional procedures

You can perform the following tasks after the 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 Boot” on page 69.
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 a compressed file that you can send to customer support.	See “Diagnostic utilities provided by NetApp” on page 138.

Installing FC HP-UX Host Utilities 4.3

Before you begin

FC HP-UX Host Utilities 4.3 provides support for HP-UX using LVM and Veritas VxVM.

Note

FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

Perl (/opt/perl/bin/perl) must be installed on the host.

Host Utilities consists of the SAN Toolkit software package.

Before you install the SAN Toolkit software package, perform the following actions:

- ◆ You must verify that your host system meets Host Utilities requirements. See one of the following references:
 - ❖ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
 - ❖ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml for the latest hardware and software requirements.
- ◆ If you currently have any previous version of the FC HP-UX Attach Kit or Host Utilities SAN Toolkit installed, you must remove it. See “[Removing FC HP-UX Host Utilities SAN Toolkit or Attach Kit](#)” on page 15.

After 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.

Step	Action
1	Use the <code>gunzip</code> command to uncompress the files.
2	Execute the <code>swinstall</code> command to install the NetApp software.

Note

You need not provide a depot path when issuing 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 appropriate 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 site at <http://now.netapp.com/NOW/cgi-bin/software> to a working directory on the HP-UX host.

Note

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

Installing the SAN Toolkit software

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

Step	Action
1	Enter the following command to download the <code>netapp_fcp_hpux_host_utilities_4.3_ia_pa.depot.gz</code> file: <pre># gunzip netapp_fcp_hpux_host_utilities_4.3_ia_pa.depot.gz</pre> Result: The system places the extracted software in the directory to which you downloaded or copied the file.

Step	Action
2	Use the <code>swinstall</code> command to install the software. You do not need to enter a path name with this command. The depot file provides the installation directory information. Attention _____ 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 downloaded the files in the <code>/tmp/hpux</code> directory on an HP-UX host called <code>machine1</code>. <pre> swinstall -s /tmp/hpux/netapp_fcp_hpux_host_utilities_4.3_ia_pa.depot NetApp_santoolkit ===== 05/20/08 14:59:12 IST BEGIN swinstall SESSION (non-interactive) (jobid=machine1-0147) * Session started for user "root@machine1". * Beginning Selection * Target connection succeeded for "machine1:/". * Source: /home/netapp_fcp_hpux_host_utilities_4.3_ia_pa.depot * Targets: machine1:/ * Software selections: NetApp_santoolkit.command_itanium,r=4.3,fr=3.3,fa=HP-UX_B.11.22_IA NetApp_santoolkit.command_parisc,r=4.3,fr=3.3,fa=HP-UX_B.11.00_32/6 4 NetApp_santoolkit.man,r=4.3 NetApp_santoolkit.support_scripts,r=4.3 * 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 machine1-0147 @ machine1:/". ===== 05/20/08 14:59:36 IST END swinstall SESSION (non-interactive) (jobid=machine1-0147) </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 the VxVM software, you must install the Symantec ASL for NetApp storage systems.

The ASL is not distributed with Host Utilities. You must download it from the Symantec Web site. For more information about using the ASL, see “[Overview of the ASL and APM](#)” on page 100.

Removing FC HP-UX Host Utilities SAN Toolkit or Attach Kit

Removing Host Utilities or Attach Kit

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

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

Uninstalling the SAN Toolkit or an Attach Kit

To uninstall the SAN Toolkit or Attach Kit software, complete the following step.

Step	Action
1	Use the <code>swremove</code> command to remove the software. ◆ To remove Host Utilities SAN Toolkit, enter the following command: <code># swremove NetApp_santoolkit</code> ◆ To remove FC HP-UX Attach Kit, enter the following command: <code># swremove santools</code> Example of removing the SAN Toolkit: The following example shows output you might see when you execute the <code>swremove</code> command to remove FC HP-UX Host Utilities SAN Toolkit from an HP-UX host called machine1. <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/santoolkit,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 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> 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 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 an overview of the tasks for configuring the storage system, and the FC cabling and switches.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“Configuring the storage system”](#) on page 20
- ◆ [“Configuring FC cabling and switches”](#) on page 22

Configuring the storage system

Before you begin

You use the WWPNs when creating igroups on the storage system.

Preparing the storage system

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

Tasks for configuring the storage system

To configure the storage system and prepare it for initiator access, you must complete the following tasks:

- ◆ Ensure that the FC service is licensed and started.
- ◆ For active/active storage system configurations, verify that the storage system is running a supported cfmode by running the `fcv show cfmode` command. The support matrix lists the supported cfmode values for your configuration.

Attention

Do not change the cfmode. If the cfmode is not set correctly, contact customer support.

Note

To determine which cfmodes are supported by FC HP-UX Host Utilities, see the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

- ◆ Create LUNs, create an igroup using the host's WWPNs, and then map LUNs to the igroup for the host. You can manage LUNs using the Web-based FilerView interface or using the Data ONTAP command line. Be sure to specify the LUN type and FC igroup type as `hpux`.
- ◆ ALUA must be enabled on the igroup if you are using an ALUA supported configuration. Check the NetApp Interoperability Matrix at

<http://now.netapp.com/NOW/products/interoperability> for the list of configurations that support ALUA.

- ◆ If you are using clustered hosts, the igroup must contain the WWPNs of all hosts in the cluster that need access to the mapped LUN.

**For detailed
configuration steps**

For detailed storage system configuration steps, see the *Block Access Management Guide* for your version of Data ONTAP.

Configuring FC cabling and switches

Fabric or direct connections

You can connect your host to the storage system by using either Fibre Channel switches or direct cables. A switched fabric provides more flexibility and enables you to connect more hosts; a direct connect requires one target port on each storage system controller for each host.

Switched fabrics

High availability: To avoid single points of failure, use two physically independent Fibre Channel switch fabrics to connect your host to an active/active storage system configuration.

Configuration and zoning details: For details about supported Fibre Channel configurations and switch zoning, see on of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Use default HBA and target topology: For a fabric-connected configuration, the default autonegotiated HBA and storage system target Fibre Channel topology works correctly.

Single initiator zoning: When using the switched fabric, you should use single initiator zoning for each of the initiator ports; each zone should have only a single initiator port.

Direct connections

High availability: To avoid single points of failure, use two Fibre Channel cables to connect your host to both controllers of an active/active storage system configuration.

Use default HBA and target topology: For a direct connected configuration, the default autonegotiated HBA and storage system's target Fibre Channel topology works correctly.

Single_image cfmodes required: A direct Fibre Channel connection to an active/active storage system configuration requires the storage system to use single_image cfmodes.

To verify cfmode, enter the following command on the storage system console:

```
fcg show cfmode
```


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 to 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 an HP-UX host”](#) on page 26
- ◆ [“Creating igroups and LUNs”](#) on page 27
- ◆ [“Configuring LUNs”](#) on page 38
- ◆ [“Displaying LUN and HBA information”](#) on page 61

Overview of adding LUNs to an HP-UX host

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 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 HP-UX LVM or VxVM.

Note

To perform some of these tasks, you must use the `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 14).

For an HP-UX host, the following `cfmodes` are the default settings:

- ◆ FAS270s: dual-fabric mode
- ◆ All other models and backup storage systems: partner.

In some circumstances, you will want to use a different `cfmode` setting. Consult with NetApp professional services if you have any questions about which mode you should use.

The standby `cfmode` setting is not supported in configurations with HP-UX hosts.

Creating igroups and LUNs

Creating and mapping LUNs on the storage system

To configure the NetApp storage controller, you must perform the following tasks:

Step	Action
1	Create an igroup on the storage system. ◆ Each igroup should map to a host and should include all the initiators used by the host to access its data.
2	Create one or more LUNs on the storage system and map LUNs to an igroup. ◆ Ensure that tLUNs 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.

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.
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 the <code>sanlun fcp show adapter -c</code> command, it displays the igroup command which contains the host's WWPNs. 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. Enter this filer command to create an initiator group for this system: <pre>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 will copy the following command line: <pre>igroup create -f -t hpux "hostA" 50060b000010a700 50060b000010a606</pre>
5	Paste the <code>igroup create</code> 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 the output of the <code>igroup show</code> command. Note If you change the <code>ostype</code> attribute to <code>hpux</code> after you have created the igroup, you must manually set the <code>vsa</code> attribute, which HP-UX sets correctly by default. For information about manually setting the <code>vsa</code> attribute, see “Changing the vsa value” on page 126. <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,

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 create` command in command-line mode. Instead of answering the prompts from the interactive script, supply the *size*, *type*, *lun_path* 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 hpx to specify the host operating system: <pre>Multiprotocol type of LUN (image/solaris/windows/hpx/aix/linux) [image]: hpx</pre>

Step	Action
4	Enter the LUN path: Attention _____ 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: lun_path <i>lun_path</i> is the path name of the LUN. Result: The name of the LUN (the LUN path) is defined. Example: The following are some examples of a <i>lun_path</i>: /vol/vol1/qtree1/lun_sdata /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: x_x <i>x_x</i> 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, which means 5 GB of raw disk. Enter LUN size: 5g
7	Enter a comment string, 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 <code>?</code> 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 <code>'?'</code> to see existing initiator group names. Name of initiator group [<i>default_igroup</i>]: new_group_name <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 <code>hpux_groupB</code> or <code>hostA</code>. Result: The program either creates a new igroup or uses the existing group that you specify. Example 1: This example uses an existing igroup. <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 <code>hpuxB</code>. <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 the setup information for the igroup. Press Enter to accept FC as the protocol: Type of initiator group <i>new_group_name</i> (FCP/iSCSI) [FCP]: Result: Choosing FC means that the initiator WWPNs in this igroup are using FC to access the LUN. Note_____ If you created an igroup as shown in “Creating an igroup” on page 27, 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 port names separated by commas 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 <code>sanlun fcp show</code> 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: <pre>OS type of initiator group "new_group_name" (default/solaris/windows/hpux/aix/linux) :hpux</pre> 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 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 the preceding 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 as n or y respectively. Do you want to create another LUN? [n]:

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#: OcABC1234ABC Share: none Multiprotocol Type: hpx Maps: hostA=3333 filerX></pre>

Configuring 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.

Note_____FC HP-UX Host Utilities 4.3 provides Veritas support with HP-UX 11iv2 but not with HP-UX 11iv3.

- ◆ Reboot the host. The host automatically discovers the new LUNs when you reboot the host.
- ◆ If rebooting the host is not a reasonable action for the sotrage system, perform the steps listed here. Performing this steps enables the host to discover LUNs without requiring a host reboot.

Step	Action
1	Log in as root on the host.

Step	Action
2	Perform an ioscan on the HP-UX host to see whether it discovers the newly created LUNs: ioscan -fn -C disk Example: In this example, the output from an ioscan command provides the following information: ◆ There are no device special files for 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 does not create device special files, use the <code>insf -e</code> command to install the device special files. 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, specify the paths that are primary; you will need this information later. 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 filer1:/vol/vol1/testlun (LUN 0) 10g (10737418240) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: filer2 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/c17t0d0 fcd1 0c up primary /dev/dsk/c14t0d0 fcd1 0c up primary /dev/dsk/c11t0d0 fcd0 0d up secondary /dev/dsk/c8t0d0 fcd0 0d</pre> Note_____ If you are using HP-UX 11iv3 February 2007, go to Step 5. If you are using HP-UX 11iv2, skip Step5. _____

Step	Action
5	Note_____ If you are using the HP-UX 11iv2, skip this step. If you are using the HP-UX 11iv3 February 2007 release, continue with this step. The HP-UX 11iv3 February 2007 release does not support ALUA. In this case, the default multipathing policy for all the disk storage is active/active. So if you are using the HP-UX 11iv3 February 2007 release, you might need to change the multipathing policy to active/passive for Data ONTAP LUNs. You can use the <code>enable_ontap_pvlinks</code> script to change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing the other vendors' LUNs. Note_____ With later versions of HP-UX 11iv3, ALUA is supported. The <code>enable_ontap_pvlinks</code> command has the following format: <pre>enable_ontap_pvlinks [set show unset]</pre> Run the following command to set the multipathing policy to active/passive for all Data ONTAP LUNs without disturbing other Vendors' LUNs: <pre>enable_ontap_pvlinks set</pre> Run the following command to display the current multipathing policy for all Data ONTAP LUNs: <pre>enable_ontap_pvlinks show</pre>

To manage the LUN using LVM, see “[Managing LUNs by using LVM](#)” on page 43.

To manage the LUN using VxVM, see “[Managing LUNs by using VxVM](#)” on page 54.

Note_____

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

Managing LUNs by using LVM

To configure LUNs as disks by using 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> 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 from the example in Step 2, 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. Do not 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</i> /dev/dsk/<i>path_disk_device</i> <i>vg_name</i> is a volume group name, with or without the <code>/dev/</code> prefix. <i>path_disk_device</i> is the path name to the disk. 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's 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 120. Example: The <code>ntap_config_paths</code> utility adds the rest of the paths using the <code>vgextend</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
	You can verify the path ordering with the following command: <pre>vgdisplay -v volume_group</pre>
7	Create one or more logical volumes in the volume group using the following command line: <pre>lvcreate [-L -l] size volume_group</pre> <i>size</i> with an uppercase <code>-L</code> option is the total size of the logical volume in megabytes. <i>size</i> with a lowercase <code>-l</code> option is the number of physical extents that you want to assign to the logical volume. The value you supply for <i>size</i> 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 following command: <pre>vgdisplay -v volume_group</pre> 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 <code>-l</code> option assigns them to the logical volume. See the man 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	To create a file system on the logical volume device special file, use the <code>newfs</code> command. <code>newfs -F vxfs /dev/volume_group/character_device</code> <i>character_device</i> is the character device listed in the output of the following command: <pre>vgdisplay -v volume_group</pre> 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	After you create the new file system, create a mountpoint for it. Example: <pre># mkdir /mnt/vg_ntap01_lv1</pre>
10	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/rlvol1 /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 PVLlinks order of the paths to the LUN. <pre># sanlun lun show -p</pre> <pre>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</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/c10t0d4</td><td>fcd1</td><td>0c</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c4t0d4</td><td>fcd0</td><td>0c</td><td>1</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c13t0d4</td><td>fcd1</td><td></td><td>0d 2</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c7t0d4</td><td>fcd0</td><td></td><td>0d 3</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/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
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/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																																																		

Managing LUNs using LVM with ALUA enabled on the igroup

If you are using HP-UX 11iv3 September 2007 release and later, you can create the LVM volume groups using the agile path names.

You must enable ALUA on the igroup before mapping LUNS.

To enable ALUA on the igroup use the following command
igroup set <igroup_name> alua on

Step	Action
1	Create a physical volume on the LUN by entering the following command: <pre><i>pvccreate /dev/rdisk/path_disk_device</i></pre> <i>path_disk_device</i> is the path to the agile device that contains the LUN. Example: The <code>pvccreate</code> command produces output similar to the following: <pre># pvccreate /dev/rdisk/disk23</pre> The physical volume <code>/dev/rdisk/disk23</code> has been successfully created. Use Steps 2, 3, and 4 from “Managing LUNs by using LVM” on page 43 about creating the device node.
2	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. <pre><i>vgcreate vg_name /dev/disk/path_disk_device</i></pre> <i>vg_name</i> is a volume group name, with or without the <code>/dev/</code> prefix. <i>path_disk_device</i> is the path name to the disk. Example: In this example, the <code>vgcreate</code> command sets up a primary path to the volume group. <pre># vgcreate /dev/vg_ntap01 /dev/disk/disk23</pre> The volume group <code>/dev/vg_ntap01</code> has been successfully created. The volume group configuration for <code>/dev/vg_ntap01</code> has been saved in <code>/etc/lvmconf/vg_ntap01.conf</code>.

Step	Action
3	Create one or more logical volumes in the volume group using the following command line: <pre>lvcreate [-L -l] size volume_group</pre> <i>size</i> with an uppercase -L option is the total size of the logical volume in megabytes. <i>size</i> 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 <i>size</i> 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 following command: <pre>vgdisplay -v volume_group</pre> 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 man 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
4	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/volume_group/character_device</code> <i>character_device</i> is the character device listed in the output of the following command: <pre>vgdisplay -v volume_group</pre> 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>
5	After you create the new file system, create a mountpoint for it. Example: <pre># mkdir /mnt/vg_ntap01_lv1</pre>
6	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																																																
7	Verify that the LUN is visible to the host by entering the following command: # sanlun lun show -p controller1:/vol/mvol2/lun2 (LUN 1) Agile Filename: /dev/rdisk/disk48 VG: /dev/vg_ntap01 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/A Multipath_Provider: HP A/A <table><tr><td>host</td><td>filer</td><td>/dev/dsk</td><td>primary</td><td>partner</td><td>HP A/A</td></tr><tr><td>path</td><td>path</td><td>filename</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>or Hardware Path</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>priority</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c15t0d1</td><td>td3</td><td>0d</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c13t0d1</td><td>td3</td><td>0c</td><td>0</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c17t0d1</td><td>td5</td><td>0c</td><td>1</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c19t0d1</td><td>td5</td><td>0d</td><td>1</td></tr></table>	host	filer	/dev/dsk	primary	partner	HP A/A	path	path	filename	host	filer	filer	state	type	or Hardware Path	HBA	port	port						priority	up	primary	/dev/dsk/c15t0d1	td3	0d	0	up	primary	/dev/dsk/c13t0d1	td3	0c	0	up	secondary	/dev/dsk/c17t0d1	td5	0c	1	up	secondary	/dev/dsk/c19t0d1	td5	0d	1
host	filer	/dev/dsk	primary	partner	HP A/A																																												
path	path	filename	host	filer	filer																																												
state	type	or Hardware Path	HBA	port	port																																												
					priority																																												
up	primary	/dev/dsk/c15t0d1	td3	0d	0																																												
up	primary	/dev/dsk/c13t0d1	td3	0c	0																																												
up	secondary	/dev/dsk/c17t0d1	td5	0c	1																																												
up	secondary	/dev/dsk/c19t0d1	td5	0d	1																																												

Note

In this case, as ALUA is enabled, you donot need to run the script enable_ntap_pvlinks as well as ntap_config_paths.

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 one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Note

FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 but not with 11iv3.

Adding a new LUN under VxVM management

To add 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 38 for details.
2	Enter the following command to enable VxVM to detect the new LUN: <code>vxctl enable</code>
3	Enter the following command to verify that VxVM detects the LUN: <code>vxdisk list</code> 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 DEVICE column. If you are using disk-based naming, then the controller or disk name (for example, c0t0d3) 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>

Step	Action
4	Enter the following command to initialize the disk for VxVM: <pre>/usr/lib/vxvm/bin/vxdisksetup -i <i>device_name</i></pre> <i>device_name</i> is the name listed for the LUN in the DEVICE 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=cdsdisk # vxdisksetup -i FAS30200_1 format=cdsdisk # vxdisk list DEVICE TYPE DISK GROUP STATUS Disk_0 auto:LVM - - LVM Disk_1 auto:LVM - - LVM Filer_0 auto:cdsdisk - - online Filer_1 auto:cdsdisk - - online</pre>
5	Enter the following command to create a disk group for the LUN: <pre>vx dg init <i>dg_name</i> <i>diskname=device_name</i></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 <code>vxdisk list</code> command output. Example: <pre># vx dg init dg1 dg1-disk1=Filer_0</pre>

Step	Action																									
6	Enter the following command to create disks or 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. <i>diskname</i> represents the device. <i>device_name</i> is the controller name listed for the LUN in the DEVICE column of the <code>vxdisk list</code> command output. Example: <pre># vxdbg -g dg1 adddisk dg1-disk2=Filer_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																						
7	Enter the following command to create a logical volume: <pre>vxassist -g dg_name make vol_name size</pre> <i>dg_name</i> is the name of the disk group that you defined in Step 6. <i>vol_name</i> is the name you assign to the logical volume. <i>size</i> is the volume size. For example, 50m (50MB). Size cannot be equal to or exceed the 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/rdsk/dg_name/vol_name</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.																									

Step	Action
9	Create a mountpoint 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>
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 <code>DMP-NODE</code>. If you are using disk device-based naming, the <code>DMP-NODE</code> is listed in the device-name format (for example, <code>c16t0d1</code>). Note that the output shows the Veritas path I/O policy as Not Applicable. The policy is single active but <code>sanlun</code> does not identify which path is currently active. <pre># sanlun lun show -p controller1:/vol/vol1/lun1 (LUN 0) DMP-NODE: (FAS30700_94) 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/PG Multipath_Provider: Veritas DMP ----- host filer path path /dev/dsk host filer filer path I/O state type filename HBA port port policy ----- up secondary /dev/dsk/c12t0d0 fcd4 0d Round-Robin up secondary /dev/dsk/c10t0d0 fcd4 0c Round-Robin up primary /dev/dsk/c6t0d0 fcd3 0d Round-Robin up primary /dev/dsk/c4t0d0 fcd3 0c Round-Robin</pre>

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.

You should use the following values for these settings:

- ◆ restore daemon interval: 60
- ◆ restore daemon policy for Veritas Versions

	Veritas 4.1	Veritas 5.0
Data ONTAP 7.0.x	check_disabled	check_disabled
Data ONTAP 7.1.x	check_disabled	check_disabled
Data ONTAP 7.2.x	check_all	check_disabled
Data ONTAP 7.3.x	check_all	check_disabled

	Veritas 4.1	Veritas 5.0
vxpfto	120 seconds	30 seconds

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 following 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 following command: <pre># vxddmpadm stop restore</pre>

Step	Action
4	Start the restore daemon with the values set to the following: ◆ 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>
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
6	For Veritas 4.1 set vxpf to 120 seconds. Changing vxpf: <pre>/opt/VRTS/bin/vxpf -t 120 -g <VGNAME></pre> For Veritas 5.0 set vxpf to 30 seconds Changing vxpf: The default value is 30 seconds, so no need of setting the timeout value.

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 the host and the 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 HBA.

For detailed information

For detailed information, see the following topics:

- ◆ [“Displaying host LUN information”](#) on page 62
- ◆ [“Displaying HBA information”](#) on page 67

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 PVLlinks or Veritas DMP) and the multipathing policy. For example, if you are using PVLlinks, 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> produces verbose output. <code>-d</code> specifies the special device file on the host (<i>host_device_filename</i>); for example, <code>-d /dev/rdisk/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 storage system. <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 all parameters.

Step	Action
	The command line options 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 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 LUNs associated with the host device /dev/rdisk/c3t4d2, enter the following command line: sanlun lun show -d /dev/rdisk/c3t4d2 ◆ To display verbose output for all LUNs currently available on the host, enter the following command line: sanlun lun show -v all ◆ To display a summary listing of all 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 through 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 (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.

Step	Action																																																																																																												
	Example using LVM: On a system using LVM, you will see output similar to the following if you execute the <code>sanlun lun show</code> command with the <code>-p</code> option before running the <code>ntap_config_paths</code> utility. In this case, LUNs are managed by LVM and PVLlinks. The PVLlinks failover priority column shows the PVLlinks order of the paths to the LUN. # sanlun lun show -p <pre>controller1:/vol/mvol1/lun1 (LUN 5) Agile Filename: /dev/rdisk/disk24 VG: /dev/legvg lg (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/P Multipath-provider: HP PV Links</pre> <table><tr><th>host</th><th>filer</th><th>/dev/dsk</th><th>Primary</th><th>partner</th><th>PVlinks</th></tr><tr><th>path</th><th>path</th><th>filename</th><th>host</th><th>filer</th><th>filer</th></tr><tr><th>state</th><th>type</th><th>or Hardware Path</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/c5t0d5</td><td>fcd5</td><td>0d</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c11t0d5</td><td>fcd5</td><td>0b</td><td>1</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c9t0d5</td><td>fcd4</td><td>0b</td><td>2</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c7t0d5</td><td>fcd4</td><td>0d</td><td>3</td></tr></table> Example using Veritas DMP: On a system using DMP, you will see output similar to the following if you execute the <code>sanlun lun show</code> command. # sanlun lun show -p <pre>controller1:/vol/vol1/lun1 (LUN 0) DMP-NODE: (FAS30700_94) lg (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/PG Multipath_Provider: Veritas DMP</pre> <table><tr><th>host</th><th>filer</th><th>/dev/dsk</th><th>primary</th><th>partner</th><th>Veritas</th></tr><tr><th>path</th><th>path</th><th>filename</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/c12t0d0</td><td>fcd4</td><td>0d</td><td>Round-Robin</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c10t0d0</td><td>fcd4</td><td>0c</td><td>Round-Robin</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c6t0d0</td><td>fcd3</td><td>0d</td><td>Round-Robin</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c4t0d0</td><td>fcd3</td><td>0c</td><td>Round-Robin</td></tr></table>	host	filer	/dev/dsk	Primary	partner	PVlinks	path	path	filename	host	filer	filer	state	type	or Hardware Path	HBA	port	port						path failover						priority	up	primary	/dev/dsk/c5t0d5	fcd5	0d	0	up	primary	/dev/dsk/c11t0d5	fcd5	0b	1	up	secondary	/dev/dsk/c9t0d5	fcd4	0b	2	up	secondary	/dev/dsk/c7t0d5	fcd4	0d	3	host	filer	/dev/dsk	primary	partner	Veritas	path	path	filename	host	filer	filer	state	type	filename	HBA	port	port						path I/O						policy	up	secondary	/dev/dsk/c12t0d0	fcd4	0d	Round-Robin	up	secondary	/dev/dsk/c10t0d0	fcd4	0c	Round-Robin	up	primary	/dev/dsk/c6t0d0	fcd3	0d	Round-Robin	up	primary	/dev/dsk/c4t0d0	fcd3	0c	Round-Robin
host	filer	/dev/dsk	Primary	partner	PVlinks																																																																																																								
path	path	filename	host	filer	filer																																																																																																								
state	type	or Hardware Path	HBA	port	port																																																																																																								
					path failover																																																																																																								
					priority																																																																																																								
up	primary	/dev/dsk/c5t0d5	fcd5	0d	0																																																																																																								
up	primary	/dev/dsk/c11t0d5	fcd5	0b	1																																																																																																								
up	secondary	/dev/dsk/c9t0d5	fcd4	0b	2																																																																																																								
up	secondary	/dev/dsk/c7t0d5	fcd4	0d	3																																																																																																								
host	filer	/dev/dsk	primary	partner	Veritas																																																																																																								
path	path	filename	host	filer	filer																																																																																																								
state	type	filename	HBA	port	port																																																																																																								
					path I/O																																																																																																								
					policy																																																																																																								
up	secondary	/dev/dsk/c12t0d0	fcd4	0d	Round-Robin																																																																																																								
up	secondary	/dev/dsk/c10t0d0	fcd4	0c	Round-Robin																																																																																																								
up	primary	/dev/dsk/c6t0d0	fcd3	0d	Round-Robin																																																																																																								
up	primary	/dev/dsk/c4t0d0	fcd3	0c	Round-Robin																																																																																																								

Step	Action
	Example using LVM with ALUA: On an HP-UX 11iv3 September 2007 release or later system using LVM with ALUA enabled on the adapter igroup, you will see following sanlun output. In this case, LUNs are managed by LVM and PVLlinks. The PVLlinks failover priority column shows the PVLlinks order of the paths to the LUN. <pre># sanlun lun show -p controller1:/vol/mvol2/lun2 (LUN 1) Agile Filename: /dev/rdisk/disk48 VG: /dev/vg01 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/A Multipath_Provider: HP A/A ----- host filer /dev/dsk primary partner HP A/A path path filename host filer filer path failover state type or Hardware Path HBA port port priority ----- up primary /dev/dsk/c15t0d1 td3 0d 0 0 up primary /dev/dsk/c13t0d1 td3 0c 0 0 up secondary /dev/dsk/c17t0d1 td5 0c 1 1 up secondary /dev/dsk/c19t0d1 td5 0d 1 1</pre>

Step	Action																																																
	Example using LVM with ALUA while leagacy mode is disabled: On a HP-UX 11iv3 September 2007 release or later system using LVM with ALUA enabled on the controller igroup and the host has legacy I/O nodes and legacy DSFs disabled, the following output is shown by sanlun: <pre># sanlun lun show -p</pre> <pre>controller1:/vol/mvol2/lun2 (LUN 1) Agile Filename: /dev/rdisk/disk48 VG: /dev/vg01 1g (1073741824) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 Multipath_Policy: A/A Multipath_Provider: HP A/A</pre> <table><tr><td>host</td><td>filer</td><td>/dev/dsk</td><td>primary</td><td>partner</td><td>HP A/A</td></tr><tr><td>path</td><td>path</td><td>filename</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>or Hardware Path</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>priority</td></tr><tr><td>up</td><td>primary</td><td>1/0/6/1/0.0x500a098296d81673.0x4001000000000000</td><td>td3</td><td>0d</td><td>0</td></tr><tr><td>up</td><td>primary</td><td>1/0/6/1/0.0x500a098196d81673.0x4001000000000000</td><td>td3</td><td>0c</td><td>0</td></tr><tr><td>up</td><td>secondary</td><td>1/0/14/1/0.0x500a098186d81673.0x4001000000000000</td><td>td5</td><td>0c</td><td>1</td></tr><tr><td>up</td><td>secondary</td><td>1/0/14/1/0.0x500a098286d81673.0x4001000000000000</td><td>td5</td><td>0d</td><td>1</td></tr></table>	host	filer	/dev/dsk	primary	partner	HP A/A	path	path	filename	host	filer	filer	state	type	or Hardware Path	HBA	port	port						priority	up	primary	1/0/6/1/0.0x500a098296d81673.0x4001000000000000	td3	0d	0	up	primary	1/0/6/1/0.0x500a098196d81673.0x4001000000000000	td3	0c	0	up	secondary	1/0/14/1/0.0x500a098186d81673.0x4001000000000000	td5	0c	1	up	secondary	1/0/14/1/0.0x500a098286d81673.0x4001000000000000	td5	0d	1
host	filer	/dev/dsk	primary	partner	HP A/A																																												
path	path	filename	host	filer	filer																																												
state	type	or Hardware Path	HBA	port	port																																												
					priority																																												
up	primary	1/0/6/1/0.0x500a098296d81673.0x4001000000000000	td3	0d	0																																												
up	primary	1/0/6/1/0.0x500a098196d81673.0x4001000000000000	td3	0c	0																																												
up	secondary	1/0/14/1/0.0x500a098186d81673.0x4001000000000000	td5	0c	1																																												
up	secondary	1/0/14/1/0.0x500a098286d81673.0x4001000000000000	td5	0d	1																																												
	Example using -v: The following is an example of output you might see when you enter the sanlun command with the -v option only. <pre># sanlun lun show -v</pre> <pre>filer: lun-pathname device filename adapter lun size lun state filerX /vol/vol1/3333 /dev/rdsk/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 HBA information

Displaying HBA information To display the 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 HBA information. This command has the following format: <pre>sanlun fcp show adapter [-c [-v] [adapter name all]]</pre> <code>-v</code> produces verbose output. <code>-c</code> gives configuration information that you can use to create igroups. <code>all</code> lists information for all FC adapters. <i>adapter name</i> lists adapter information for the adapter you specify. Example: The following example shows the type of output you see when you use different options of the <code>sanlun fcp show adapter</code> command. <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 boot is using a SAN-attached disk, such as a NetApp LUN, as a boot device for a host. SAN boot provides the following advantages:

- ◆ You can remove the hard drives from your servers and use the SAN for your booting 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 boot with LVM and another section on using it with Veritas.

For detailed information about supported components of SAN boot configurations, see one of the following references:

- ◆ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ◆ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.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 boot procedures](#)” on page 70
- ◆ “[Creating the boot LUN using LVM \(LVM only\)](#)” on page 73
- ◆ “[Migrating an OS disk using MirrorDisk/UX on PA-RISC systems \(LVM only\)](#)” on page 79
- ◆ “[Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems \(LVM only\)](#)” on page 87
- ◆ “[Testing the alternate boot path for PA-RISC systems \(LVM only\)](#)” on page 93
- ◆ “[Troubleshooting sanboot configurations with PA-RISC systems \(LVM only\)](#)” on page 94
- ◆ “[Creating the boot LUN using vxcp_lvmroot \(Veritas only\)](#)” on page 96
- ◆ “[Creating the boot LUN using a cold install from the HP-UX DVD \(Veritas only\)](#)” on page 97

Overview of SAN boot 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. See the HP® documentation for your operating system (OS) and server for more details.
- ◆ 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 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 any one LVM-using method or any one VxVM-using method .

Note

FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

Using LVM to configure the HP-UX OS on the boot LUN:

- ❖ Migrate an existing OS disk to the boot LUN by using MirrorDisk/UX. See “[Migrating an OS disk using MirrorDisk/UX on PA-RISC systems \(LVM only\)](#)” on page 79.
- ❖ Install the OS by using an Ignite Server or the HP installation CD-ROM. See “[Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems \(LVM only\)](#)” on page 87.

Using Veritas to configure the HP-UX OS on the boot LUN:

- ❖ Create a Veritas-controlled SAN Boot LUN on HP-UX. See “[Creating the boot LUN using vxcp_lvmroot \(Veritas only\)](#)” on page 96.
- ❖ Perform a cold install directly from the HP-UX Core OS DVD. See “[Creating the boot LUN using a cold install from the HP-UX DVD \(Veritas only\)](#)” on page 97.
- ◆ 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 the HP documentation for your OS and server. Itanium systems divides the boot device into three part:

- ◆ s0 represents the whole disk, for example c0t0d0s0.
- ◆ s1 represents the Extensible Firmware Interface (EFI) 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.

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

```

controller1:/vol/vol1/bootlun (LUN 101) VG: /dev/vg00
    18g (19327352832)    lun state: GOOD
Filer_CF_State: Cluster Enabled  Filer Partner: controller2
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	primary	/dev/dsk/c7t12d5s2	td0	4a	0
up	primary	/dev/dsk/c13t12d5s2	td1	5a	1
up	secondary	/dev/dsk/c11t12d5s2	td1		5b
up	secondary	/dev/dsk/c5t12d5s2	td0		4b
					3

Creating the boot LUN using LVM (LVM only)

- Procedure overview**
- To create a bootable SAN LUN using LVM, you complete the following tasks:
1.

Create the LUN on which you want to install the HP-UX OS. This includes determining the boot LUN's size, mapping it to the host's igroup, and discovering the new LUN on the host.
2.

Configure the HP-UX OS on the boot LUN using one of the following methods:

❖

Migrate an existing OS disk to the boot LUN by using MirrorDisk/UX.

❖

Install the OS by using an Ignite Server or the HP installation CD-ROM.
3.

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 Physical volumes shows the paths to the system disk. Example: The following example shows the path to the system disk. --- Physical volumes --- PV Name /dev/dsk/c1t15d0 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 the 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 27. 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 75.

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 <i>size</i> -t <i>ostype</i> <i>lun-path</i></pre> <i>size</i> is the size of the LUN, which you obtained in “Determining the size of the boot LUN” on page 73. 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), or t (terabytes). <i>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 <i>lun-path</i> <i>igroup-name</i> <i>LUN-ID</i></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 LUNs yet. - The LUN is visible to the host and the 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>
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 does not 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 “ Migrating an OS disk using MirrorDisk/UX on PA-RISC systems (LVM only) ” on page 79.
	You want to reinstall the OS using an Ignite Server or an installation CD-ROM	Proceed to “ Installing the OS using an Ignite Server or installation CD-ROM on PA-RISC systems (LVM only) ” on page 87.

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

Mirroring an existing system disk

If you have created a new LUN for the boot device, and you want to migrate the existing system disk to this 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. 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><th>host</th><th>filer</th><th></th><th></th><th>primary</th><th>partner</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><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																																							

Step	Action
2	The HP-UX 11iv3 February 2007 release, does not support ALUA. In this case, the default multipathing policy for all the disk storage is active/active. So if you are using the HP-UX 11iv3 February 2007 release, you need to change the multipathing policy to active/passive for Data ONTAP LUNs. Note_____ If you are using HP-UX 11iv2, you do not need to use this script. You can use the <code>enable_ontap_pvlinks</code> script to change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing the other vendors' LUNs. Note_____ With later versions of HP-UX 11iv3, ALUA is supported. The <code>enable_ontap_pvlinks</code> command has the following format: <pre>enable_ontap_pvlinks [set show unset]</pre> Run the following command to set the multipathing policy to active/passive for all Data ONTAP LUNs without disturbing other vendors' LUNs: <pre>enable_ontap_pvlinks set</pre> Run the following command to display the current multipathing policy for all Data ONTAP LUNs: <pre>enable_ontap_pvlinks show</pre>
3	Create a bootable physical volume on the LUN by entering the following command: <pre>pvccreate -fB /dev/rdisk/ntap_device</pre> <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>pvccreate -fB /dev/rdisk/c24t0d0</code>

Step	Action
4	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>
5	Add the remaining paths to the LUN to the root volume group by using <code>ntap_config_paths</code> utility, which is in the <code>/opt/NetApp/santools/bin</code> directory. <pre>ntap_config_paths vg00</pre>
6	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>.
7	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
8	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. Mirror the logical volumes to the boot LUN. vgdisplay -v vg00 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>
9	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: lvextend -m 1 <i>lv_path</i> /dev/dsk/<i>ntap_device</i> <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>
10	Update the Boot Device Reserved Area (BDRA) by entering the following command: lvlnboot -R

Step	Action
11	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 <code>SDUX</code> utility. Use the path to the block device node. 1 /dev/dsk/ntap_device <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 79. 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 the host is /dev/dsk/c24t0d0. This is the path to which you mirrored logical volumes and copied HP-UX boot utilities in “Mirroring an existing system disk” on page 79. 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 colspan="4">-----</td><td colspan="2">-----</td></tr><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="4">-----</td><td colspan="2">-----</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</pre> <table><tr><th>Class</th><th>I</th><th>H/W Path</th><th>Driver</th><th>S/W State</th><th>H/W Type</th><th>Description</th></tr><tr><td colspan="7">=====</td></tr><tr><td>disk</td><td>0</td><td>0/0/1/1.15.0</td><td>sdisk</td><td>CLAIMED</td><td>DEVICE</td><td>HP 18.2GMAN3184MC</td></tr><tr><td></td><td></td><td></td><td>/dev/dsk/c1t15d0</td><td>/dev/rdsk/c1t15d0</td><td></td><td></td></tr><tr><td>disk</td><td>14</td><td>0/4/0/0.1.4.0.0.0.0</td><td>sdisk</td><td>CLAIMED</td><td>DEVICE</td><td>NETAPP LUN</td></tr><tr><td></td><td></td><td></td><td>/dev/dsk/c24t0d0</td><td>/dev/rdsk/c24t0d0</td><td></td><td></td></tr><tr><td>disk</td><td>28</td><td>0/4/0/0.1.7.0.0.0.0</td><td>sdisk</td><td>CLAIMED</td><td>DEVICE</td><td>NETAPP LUN</td></tr><tr><td></td><td></td><td></td><td>/dev/dsk/c26t0d0</td><td>/dev/rdsk/c26t0d0</td><td></td><td></td></tr><tr><td>disk</td><td>37</td><td>0/6/2/0.1.4.0.0.0.0</td><td>sdisk</td><td>CLAIMED</td><td>DEVICE</td><td>NETAPP LUN</td></tr><tr><td></td><td></td><td></td><td>/dev/dsk/c46t0d0</td><td>/dev/rdsk/c46t0d0</td><td></td><td></td></tr><tr><td>disk</td><td>46</td><td>0/6/2/0.1.7.0.0.0.0</td><td>sdisk</td><td>CLAIMED</td><td>DEVICE</td><td>NETAPP LUN</td></tr><tr><td></td><td></td><td></td><td>/dev/dsk/c44t0d0</td><td>/dev/rdsk/c44t0d0</td><td></td><td></td></tr></table>	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	0/4/0/0.1.4.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN				/dev/dsk/c24t0d0	/dev/rdsk/c24t0d0			disk	28	0/4/0/0.1.7.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN				/dev/dsk/c26t0d0	/dev/rdsk/c26t0d0			disk	37	0/6/2/0.1.4.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN				/dev/dsk/c46t0d0	/dev/rdsk/c46t0d0			disk	46	0/6/2/0.1.7.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN				/dev/dsk/c44t0d0	/dev/rdsk/c44t0d0		
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	0/4/0/0.1.4.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN																																																																															
			/dev/dsk/c24t0d0	/dev/rdsk/c24t0d0																																																																																	
disk	28	0/4/0/0.1.7.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN																																																																															
			/dev/dsk/c26t0d0	/dev/rdsk/c26t0d0																																																																																	
disk	37	0/6/2/0.1.4.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN																																																																															
			/dev/dsk/c46t0d0	/dev/rdsk/c46t0d0																																																																																	
disk	46	0/6/2/0.1.7.0.0.0.0	sdisk	CLAIMED	DEVICE	NETAPP LUN																																																																															
			/dev/dsk/c44t0d0	/dev/rdsk/c44t0d0																																																																																	
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, 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>																																																																																				
4	Set the host’s alternate boot path by entering the following command: <pre>setboot -a sec_hw_path_to_lun</pre> <i>sec_hw_path_to_lun</i> is the secondary hardware path to the boot LUN from the host’s HBA2, 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>																																																																																				

Step	Action	
5	Test the primary boot path by shutting down and rebooting the system. <code>shutdown -r -y 0</code> 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. _____	
6	If...	Then...
	You cannot reboot the system.	Proceed to “ Troubleshooting sanboot configurations with PA-RISC systems (LVM only) ” on page 94.
	You can reboot the system	Proceed to “ Testing the alternate boot path for PA-RISC systems (LVM only) ” on page 93.

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

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 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.

Step	Action
4	The HP-UX 11iv3 February 2007 release, does not support ALUA. In this case, the default multipathing policy for all the disk storage is active/active. So if you are using the HP-UX 11iv3 February 2007 release, you need to change the multipathing policy to active/passive for Data ONTAP LUNs. You can use the <code>enable_ontap_pvlinks</code> script to change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing other vendors' LUNs. Note_____ With later versions of HP-UX 11iv3, ALUA is supported. _____ The <code>enable_ontap_pvlinks</code> command has the following format: <pre>enable_ontap_pvlinks [set show unset]</pre> Run the following command to set the multipathing policy to active/passive for all Data ONTAP LUNs without disturbing other Vendors' LUNs: <pre>enable_ontap_pvlinks set</pre> Run the following command to display the current multipathing policy for all Data ONTAP LUNs: <pre>enable_ontap_pvlinks show</pre> Note_____ If you are using HP-UX 11iv2, you do not need to use this script. _____

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

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 <i>root disk</i> 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 87. In the Ignite-UX Interface, you define the <i>root disk</i> in the Basic tab. When you define the primary path to the boot LUN as the <i>root disk</i>, 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 FC HP-UX Host Utilities software on the boot LUN.
4	Add the remaining paths to the LUN to the root volume group by using the NetApp’s <code>ntap_config_paths</code> command, which is in the <code>/opt/NetApp/santools/bin</code> directory: <code>ntap_config_paths vg00</code>
5	Set the host’s alternate boot path by entering the following command: <code>setboot -a sec_hw_path_to_lun</code> <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 87. 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: <code>shutdown -r -y 0</code> Note_____ If Autoboot is disabled (set to OFF) for your system, you will have to boot the system manually. _____

7	If...	Then...
	You cannot reboot the system.	Proceed to “ Troubleshooting sanboot configurations with PA-RISC systems (LVM only) ” on page 94.
	You can reboot the system	Proceed to “ Testing the alternate boot path for PA-RISC systems (LVM only) ” on page 93.

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

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 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 BCH, boot from the primary boot path by entering the following command: <code>bo pri</code>

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

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 BCH menu, reset the system to return to the BCH prompt.
2	Boot from the internal system disk by entering the following command: bo path <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: ◆ sanlun lun show -p—Shows the primary and secondary device paths to the boot LUN. ◆ ioscan -funC disk—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>

Creating the boot LUN using vxcp_lvmroot (Veritas only)

Prerequisites specific to Veritas

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

Note _____
FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

- ◆ Veritas and the NetApp ASL should already be installed on the host. See “[Overview of the ASL and APM](#)” on page 100 for information about 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 25 for information about 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 25 for more information.

Using vxcp_lvmroot to create the boot LUN

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

Step	Action
1	Copy the existing LVM root disk onto the NetApp LUN by following the vxcp_lvmroot procedure contained in the <i>Veritas Volume Manager 5.0 Administrator's Guide for HP-UX</i> .
2	Verify that the vxpf to and the restore daemon settings are tuned according to the NetApp recommendations.
3	Configure the host's primary and alternate boot path settings using the setboot command.

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

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 25 for information about creating and working with LUNs.

Note _____
FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

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 SAN boot LUN controlled by Veritas 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 <i>root disk</i> , Ignite-UX displays only <i>one</i> 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 to the NetApp LUN.
3	Install the NetApp ASL on the host. For more information, see “ Overview of the ASL and APM ” on page 100.
4	Install the SAN Toolkit on the host. For more information, see “ Installing Host Utilities ” on page 7.
5	Make the other paths to the NetApp LUN visible to the host.

Step	Action
6	Run the appropriate HP-UX and VxVM commands to discover and use the new paths. You can use <code>ioscan</code> , <code>insf -e</code> , and <code>vxdctl enable</code> .
7	Verify that the <code>vxpfto</code> and restore daemon settings are tuned according to the NetApp recommendations.
8	Configure the host's primary and alternate boot path settings using the <code>setboot</code> command.

Note

FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

About this chapter

If you are using Veritas Storage Foundation for multipathing, you must install and configure the Symantec ASL and APM for NetApp storage systems. This chapter describes how to install and configure the ASL and APM, and how these components affect multipathing and failover.

Note

To determine which version of the Veritas Storage Foundation and ASL and APM to use, check the NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>, or the NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml.

Topics in this chapter

Topics in this appendix include the following:

- ◆ “[Overview of the ASL and APM](#)” on page 100
- ◆ “[Installing the ASL and APM](#)” on page 102
- ◆ “[Removing the ASL and APM](#)” on page 105
- ◆ “[Obtaining multipathing information](#)” on page 109
- ◆ “[Understanding Array Support Library error message types](#)” on page 114

Overview of the ASL and APM

What the ASL is

The ASL is a library qualified by NetApp 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 just a bunch of disks (JBOD) 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 JBOD. 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 Storage Foundation Real Application Clusters (RAC), 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 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.

What the Array Policy Module is

The Array Policy Module (APM) is a kernel module that defines I/O error handling, failover path selection, and other failover behavior for a specific array. The NetApp APM is customized to optimize I/O error handling and failover path selection for the NetApp environment. When the ASL detects the NetApp storage array, it sets the array type to A/P-C-NETAPP. This instructs Veritas software to use the NetApp APM to handle I/O and path failures for the NetApp storage array.

Installing the ASL and APM

ASL and APM installation overview

Installing the ASL and APM involves the following tasks:

- ◆ Determining the ASL and APM version
- ◆ Obtaining the ASL and APM
- ◆ Verifying that your configuration meets system requirements

See one of the following reference for current information about the system requirements:

- ❖ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
- ❖ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Determining the ASL version

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

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

```
# vxddladm listversion
LIB_NAME                               ASL_VERSION           Min. VXVM version
=====
=
libvxxp256.sl                         vm-4.0-rev-1          4.1
libvxveritas.sl                       vm-4.0-rev-1          4.1
libvxva.sl                            vm-4.0-rev-1          4.1
libvxstorcomp.sl                     vm-4.0-rev-1          4.1
libvxshark.sl                        vm-4.0-rev-1          4.1
libvxnetapp.sl                       vm-4.1-rev-3          4.1
libvxnec.sl                          vm-4.0-rev-1          4.1
libvxhpalua.sl                       vm-4.0-rev-3          4.1
libvxhitachi.sl                     vm-4.0-rev-1          4.1
libvxhdsalua.sl                     vm-4.1-rev-1          4.1
libvxhds.sl                          vm-4.0-rev-1          4.1
libvxfujitsu.sl                     vm-4.0-rev-1          4.1
libvxfc60.sl                        vm-4.0-rev-1          4.1
libvxemc.sl                         vm-4.1-rev-2          4.1
libvxeccs.sl                        vm-4.0-rev-1          4.1
```

libvxdgc.sl	vm-4.0-rev-1	4.1
libvxcscovrts.sl	vm-4.1-rev-1	4.1
libvxautoraid.sl	vm-4.0-rev-1	4.1

Obtaining the ASL and APM software

The ASL and APM are not distributed with Host Utilities. You can obtain the ASL and APM from the Symantec Web site. Download the software before you start the installation.

For a direct link to the ASL and APM on the Symantec Web site, see the FC 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 and APM TechNote, which you can review at the Symantec Web site. The TechNote contains Symantec's instructions for installing the ASL and APM.

Installing the ASL and APM software

To install a fresh version of the ASL and APM 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 the following command: vxddladm rmjbod vid=NETAPP
3	Install the ASL and APM according to the directions provided in the ASL and APM TechNote on the Symantec Web site. You should have your LUNs set up before you install the ASL and APM. For information about setting up LUNs, see “Creating, Configuring, and Managing LUNs” on page 25.

Step	Action																																		
4	If your host is connected to NetApp storage, you can verify the installation by entering the following command and then locating the NetApp enclosure type in the output of the vxddmpadm listenclosure all 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 vxddladm listsupport all command. <pre># vxddladm listsupport all</pre> <table><thead><tr><th>LIBNAME</th><th>VID</th></tr></thead><tbody><tr><td colspan="2">=====</td></tr><tr><td>libvxautoraid.sl</td><td>HP</td></tr><tr><td>libvxcscovrts.sl</td><td>CSCOVRTS</td></tr><tr><td>libvxCLARiion.sl</td><td>DGC</td></tr><tr><td>libvxemc.sl</td><td>EMC</td></tr><tr><td>libvxfc60.sl</td><td>HP</td></tr><tr><td>libvxva.sl</td><td>HP</td></tr><tr><td>libvxhdsalua.sl</td><td>HITACHI</td></tr><tr><td>libvxhpalua.sl</td><td>HP, COMPAQ</td></tr><tr><td>libvxshark.sl</td><td>IBM</td></tr><tr><td>libvxhds9980.sl</td><td>HITACHI</td></tr><tr><td>libvxhdsusp.sl</td><td>HITACHI</td></tr><tr><td>libvxibmids8k.sl</td><td>IBM</td></tr><tr><td>libvxxp1281024.sl</td><td>HP</td></tr><tr><td>libvxxp12k.sl</td><td>HP</td></tr><tr><td>libvxnetapp.sl</td><td>NETAPP</td></tr></tbody></table>	LIBNAME	VID	=====		libvxautoraid.sl	HP	libvxcscovrts.sl	CSCOVRTS	libvxCLARiion.sl	DGC	libvxemc.sl	EMC	libvxfc60.sl	HP	libvxva.sl	HP	libvxhdsalua.sl	HITACHI	libvxhpalua.sl	HP, COMPAQ	libvxshark.sl	IBM	libvxhds9980.sl	HITACHI	libvxhdsusp.sl	HITACHI	libvxibmids8k.sl	IBM	libvxxp1281024.sl	HP	libvxxp12k.sl	HP	libvxnetapp.sl	NETAPP
LIBNAME	VID																																		
=====																																			
libvxautoraid.sl	HP																																		
libvxcscovrts.sl	CSCOVRTS																																		
libvxCLARiion.sl	DGC																																		
libvxemc.sl	EMC																																		
libvxfc60.sl	HP																																		
libvxva.sl	HP																																		
libvxhdsalua.sl	HITACHI																																		
libvxhpalua.sl	HP, COMPAQ																																		
libvxshark.sl	IBM																																		
libvxhds9980.sl	HITACHI																																		
libvxhdsusp.sl	HITACHI																																		
libvxibmids8k.sl	IBM																																		
libvxxp1281024.sl	HP																																		
libvxxp12k.sl	HP																																		
libvxnetapp.sl	NETAPP																																		

Removing the ASL and APM

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 and APM support. This is also true for arrays or disks in the OTHER_DISKS category. OTHER_DISKS are local non-FCP-attached disks.

At any time, only one version of an ASL and APM that supports Data ONTAP storage systems can be installed on the host. You cannot have multiple versions of the NetApp ASL and APM installed simultaneously.

Removing the ASL and APM

To remove the ASL and APM, complete the following steps.

Step	Action
1	Log in to the host as the root user.
2	Stop all I/O to NetApp LUNs. Note_____ In a Storage Foundation RAC, you must also stop clustering on a node before you remove the ASL and APM. _____

Step	Action
3	<pre> # swremove VRTSNTAPapm ===== 05/20/08 18:28:17 IST BEGIN swremove SESSION (non-interactive) (jobid=hpx_19-0149) * Session started for user "root@hpx_19". * Beginning Selection * Target connection succeeded for "hpx_19:/". * Software selections: VRTSNTAPapm.APM_FILES,l=/,r=5.0,v=VERITAS,fr=5.0,fa=HP- UX_B.11.23_PA * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "hpx_19:/". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "hpx_19:/". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log hpx_19-0149 @ hpx_19:/". ===== 05/20/08 18:28:35 IST END swremove SESSION (non-interactive) (jobid=hpx_19-0149) # swremove VRTSNTAPas1 ===== 05/20/08 18:29:01 IST BEGIN swremove SESSION (non-interactive) (jobid=hpx_19-0150) * Session started for user "root@hpx_19". * Beginning Selection * Target connection succeeded for "hpx_19:/". * Software selections: VRTSNTAPas1.ASL_FILES,l=/,r=5.0,a=HP- UX_B.11.23_IA/PA,v=VERITAS,fr=5.0,fa=HP-UX_B.11.23_PA * Selection succeeded. </pre>

Step	Action
3	<pre> # swremove VRTSNTAPapm ===== 05/20/08 18:28:17 IST BEGIN swremove SESSION (non-interactive) (jobid=hpx_19-0149) * Session started for user "root@hpx_19". * Beginning Selection * Target connection succeeded for "hpx_19:/". * Software selections: VRTSNTAPapm.APM_FILES,l=/,r=5.0,v=VERITAS,fr=5.0,fa=HP- UX_B.11.23_PA * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "hpx_19:/". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "hpx_19:/". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log hpx_19-0149 @ hpx_19:/". ===== 05/20/08 18:28:35 IST END swremove SESSION (non-interactive) (jobid=hpx_19-0149) # swremove VRTSNTAPasl ===== 05/20/08 18:29:01 IST BEGIN swremove SESSION (non-interactive) (jobid=hpx_19-0150) * Session started for user "root@hpx_19". * Beginning Selection * Target connection succeeded for "hpx_19:/". * Software selections: VRTSNTAPasl.ASL_FILES,l=/,r=5.0,a=HP- UX_B.11.23_IA/PA,v=VERITAS,fr=5.0,fa=HP-UX_B.11.23_PA * Selection succeeded. </pre>

Step	Action
	<pre> * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "hpux_19:". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "hpux_19:". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log hpux_19-0150 @ hpux_19:". ===== 05/20/08 18:29:19 IST END swremove SESSION (non-interactive) (jobid=hpux_19-0150) </pre>

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 Active/Passive Group (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 FC 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 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

FC 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.

FC 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 Multipath_Policy as A/PG-C for clustered

storage and A/A for non-clustered (single storage system) configurations. Veritas DMP is the multipath-provider.

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

Example: The following example shows partial output of the sanlun command when a LUN is controlled by VxVM and the host is attached to a NetApp cluster.

```
controller1:/vol/vol1/vxvm_data1 (LUN 1) DMP-NODE: (F880_1)
25m (26214400) lun state: GOOD
Filer_CF_State: Cluster Enabled Filer Partner: controller2
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: vxddmpadm getsubpaths dmpnodename=device <i>device</i> is the name listed under the output of the vxdisk list command. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths dmpnodename=FAS9800_0 NAME STATE[A] PATH-TYPE[M] CTRL-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 -</pre>
3	Enter the following command to obtain path information for a HBA: vxddmpadm getsubpaths ctrl=controller_name <i>controller_name</i> is the controller displayed under CTRL-NAME in the output of the vxddmpadm getsubpaths dmpnodename command you entered in Step 2. Result: The output displays information about the paths to the filer storage (whether the path is a primary or secondary path). The output also lists the filer that the device is mapped to. Example: The following is the type of output you should see. <pre># vxddmpadm getsubpaths ctrl=c6 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 -</pre>

Using sanlun to obtain multipathing information

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

Step	Action																																																
1	On the host command line, enter the following command: <pre>sanlun lun show -p all</pre> Result: sanlun displays path information for each LUN. For LUNs managed by VxVM, the output displays the Veritas multipathing policy (A/A or A/PG-C array types) and the path I/O policy. For LUNs managed by LVM, the ouput 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 controller1:/vol/vol1/vxvm_data1 (LUN 1) DMP-NODE: (F880_1) 25m (26214400) lun state: GOOD Filer_CF_State: Cluster Enabled Filer Partner: controller2 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><th>path I/O</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</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><td>Round-Robin</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c16t0d1</td><td>td1</td><td></td><td>5b</td><td>Round-Robin</td></tr><tr><td>up</td><td>primary</td><td>/dev/dsk/c19t0d1</td><td>td0</td><td>4a</td><td></td><td>Round-Robin</td></tr><tr><td>up</td><td>secondary</td><td>/dev/dsk/c18t0d1</td><td>td0</td><td></td><td>4b</td><td>Round-Robin</td></tr></table>	host	filer		primary	partner	Veritas	path	path	/dev/dsk	host	filer	filer	path I/O	state	type	filename	HBA	port	port	policy	up	primary	/dev/dsk/c17t0d1	td1	5a		Round-Robin	up	secondary	/dev/dsk/c16t0d1	td1		5b	Round-Robin	up	primary	/dev/dsk/c19t0d1	td0	4a		Round-Robin	up	secondary	/dev/dsk/c18t0d1	td0		4b	Round-Robin
host	filer		primary	partner	Veritas																																												
path	path	/dev/dsk	host	filer	filer	path I/O																																											
state	type	filename	HBA	port	port	policy																																											
up	primary	/dev/dsk/c17t0d1	td1	5a		Round-Robin																																											
up	secondary	/dev/dsk/c16t0d1	td1		5b	Round-Robin																																											
up	primary	/dev/dsk/c19t0d1	td0	4a		Round-Robin																																											
up	secondary	/dev/dsk/c18t0d1	td0		4b	Round-Robin																																											

Understanding Array Support Library error message types

Understanding error message severity The ASL works silently and continuously 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 the 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.	

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 116
- ◆ [“Correcting problems with Veritas”](#) on page 119
- ◆ [“Correcting volume group multipathing with `ntap\_config\_paths`”](#) on page 120
- ◆ [“Handling a takeover and giveback”](#) on page 122
- ◆ [“Recovering from a host reboot during a takeover”](#) on page 123
- ◆ [“Enabling the host to identify more than eight LUNs per target HBA”](#) on page 126
- ◆ [“Commands you can use to troubleshoot problems”](#) on page 132
- ◆ [“Diagnostic utilities provided by NetApp”](#) on page 138
- ◆ [“General troubleshooting”](#) on page 143

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 FC cfmodes setting is correct and supported by your configuration. To determine which cfmodes are supported, see one of the following references:
 - ❖ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
 - ❖ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Note

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

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

Note

Keep in mind that if an HP-UX FC driver is not in use, it logs off the target port. For this reason, the `fcp 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 LUNs, the following is true:
 - ❖ The `ostype` attribute is `hpux`
 - ❖ LUNs are mapped to host igroups.
 - ❖ LUNs 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 LUNs it should.
- ◆ The software states reported by the `ioscan` command for 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 uncertain 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

An uncertain 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 that the host sees from the storage system. These situations include the following:

- ◆ 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 123)
- ◆ If you change any of the following settings on the storage system:
 - ❖ FC cfmode (do not do this unless directed to do so by customersupport)
 - ❖ igroup Volume Set Addressing (VSA) setting

If you are using LVM, see [“Correcting volume group multipathing with ntap_config_paths”](#) on page 120.

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 NetApp directs you to so.)
- ◆ 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 126.)

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.

Note

FC HP-UX Host Utilities 4.3 provides Veritas support only with HP-UX 11iv2 and not with 11iv3.

If this happens, perform the following checks:

- ◆ Verify that the `vxpfto` setting is set to the values specified in this document.
- ◆ 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 58.
- ◆ Verify the storage system is configured with a supported Data ONTAP cfmode. See one of the following references:
 - ❖ NetApp Interoperability Matrix at <http://now.netapp.com/NOW/products/interoperability>
 - ❖ NetApp SAN Support Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/index.shtml

Enclosure-based naming is not working correctly when using Veritas

When you are using the enclosure-based naming on your system, it might not function properly in some cases.

Note

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

If this happens, verify the following:

- ◆ The ASL is loaded correctly. See “[Overview of the ASL and APM](#)” on page 100.
- ◆ 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 are 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 the LVM’s 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

controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
2g (2147483648)) lun state: GOOD
Filer_CF_State: Cluster Enabled  Filer Partner: controller2
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 LUNs owned by the failed storage system.

If you have a storage system cluster, the host can use the secondary paths to access 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 LUNs. If a reboot occurs, see the section “[Recovering from a host reboot during a takeover](#)” on page 123 to determine 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 S/W 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.
 - ❖ 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 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:
- ◆ The storage system that is down loses the primary paths to LUNs it owns. It may lose its secondary paths as well.
 - ◆ The volume groups deactivate the primary paths.
 - ◆ The volume groups remain active, but are aware of the secondary paths only. These are the paths from the active storage system to LUNs on the failed storage system.
 - ◆ 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 122 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 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: <code>cf giveback</code>

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/clt15d0 /dev/rdisk/clt15d0 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>																																										
4	On the host, enter the following command to confirm that the host detects all paths to 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 Filer Partner: controller2 Multipath_Policy: A/P Multipath-provider: HP PVlinks -----</pre> <table><thead><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></thead><tbody><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></tbody></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 Filer Partner: controller2 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 eight 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 the `vsa` attribute off.

Changing the vsa value

To manually turn the `vsa` attribute 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	On the 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 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 <i>filer1 name</i> sanlun lun show -p <i>filer2 name</i></pre> <i>filer1 name, filer2 name</i> are name of the storage systems.
2	On the host Close the logical volumes. If a logical volume contains a file system, unmount the file system.

Step	Action
3	On the host Deactivate each volume group by entering the following command: <pre>vgchange -a n /dev/vol_group_name</pre> <i>vol_group_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	On the host Remove the volume group entry from the /etc/lvmtab and associated device files by using the <code>vgexport</code> command: <pre>vgexport /dev/ntap01</pre>
5	On the storage system Go into the diagnostic mode by entering the following command: <pre>priv set diag</pre> Example: When you start the 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	On the storage system Enter the following command: <pre>igroup set igroup vsa value</pre> <i>igroup</i> is the name of the igroup for which you want to change the setting. <i>value</i> is either “on” or “off”. For HP-UX systems, you must always set this value to “on.” Result: The <i>vsa</i> value of the specific igroup is set to “on” or “off”. 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 filerX to “on”: <pre>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	On the 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 data-bbox="494 618 1228 1216">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	On the 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 show up as NO_HW in the <code>ioscan</code> output.
9	On the 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	On the 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	On the host Re-add the volume group entries to /etc/lvmtab and the associated device files to the system. For each volume group, perform the following tasks: - Create a new directory with the <code>mkdir</code> command: <pre># mkdir /dev/<vg_name></pre> <i>vg-name</i> is the volume group name. - 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> - 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	On the 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	On the host Run the <code>ntap_config_paths</code> utility to configure multipathing for the newly imported volume groups: <pre># ntap_config_paths</pre>
14	On the host Verify that the paths are now set up correctly: <pre># sanlun lun show all -p</pre>
15	Repeat Step 1 to Step 14 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 `ioscan` command
- ◆ The NetApp `sanlun` utility
- ◆ The NetApp `enable_ontap_pvlinks` script

Using the HP-UX `ioscan` command

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

Example of `ioscan` output

The following example shows the type of output from the `ioscan` command you will 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 sanlun utility

You can use the `sanlun` utility to perform the following tasks:

- ◆ Check the status of the following:
 - ❖ 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 perform the following actions:

- ◆ 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]

controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
      2g (2147483648))          lun state: GOOD
Filer_CF_State: Cluster Enabled  Filer Partner: controller2
Multipath_Policy: A/P
Multipath-provider: HP PVlinks
-----
host      filer
path      path          /dev/dsk      host filer      filer  PVlinks
state     type          filename      HBA  port      port  path failover
-----
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 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` shows 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

controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
      2g (2147483648))          lun state: GOOD
Filer_CF_State: Cluster Enabled  Filer Partner: controller2
Multipath_Policy: A/P
Multipath-provider: HP PVLlinks
-----
host      filer                                primary partner PVLlinks
path      path          /dev/dsk          host filer   filer   path failover
state     type          filename         HBA  port    port    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.

Replacing a LUN on HP-UX 11iv3

If you replace a old LUN device with a new or another LUN in the given I-T-L (Initiator Target LUN) nexus, then the LUN has to go through a manual replacement process.

To replace the LUN, perform the following steps.

Step	Action
1	Stop all the I/O on the LUN device which needs to be replaced.
2	Close the LUN device.
3	Unmap the LUN device in the controller from the host.
4	Map a new LUN device in the controller at the same LUN ID, where the LUN is unmapped in the previous step.
5	Perform an ioscan on the host.

Step	Action
6	During the ioscan, you will see the following message in the host syslog (/var/adm/syslog/syslog.log). <pre>vmunix: class : lunpath, instance 37 vmunix: Evpd inquiry page 83h/80h failed or the current page 83h/80h data do not match the previous known page 83h/80h data on LUN id 0x0 probed beneath the target path (class = tgtpath, instance = 5) The lun path is (class = lunpath, instance 37).Run 'scsimgr replace_wwid' command to validate the change vmunix: vmunix: An attempt to probe existing LUN id 0x4001000000000000 failed with errno of 14.</pre>
7	Use ioscan -fnNC lunpath grep NO_HW to find the paths that turned to NO_HW due to this replacement. Enter scsimgr replace_wwid -H <lunpath> on the paths to manually validate the change in wwid of the LUN. Example: <pre># scsimgr replace_wwid -H 0/3/0/0/0/1.0x500a098197395e38.0x4001000000000000 scsimgr:WARNING: Performing replace_wwid on the resource may have some impact on system operation. Do you really want to replace? (y/[n])? y Binding of LUN path 0/3/0/0/0/1.0x500a098197395e38.0x4001000000000000 with new LUN validated successfully</pre> See scsimgr man page for more information.

Example of sanlun output after a 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 following:

- ◆ 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

controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/vg_ontap01
  2g (2147483648)   lun state: GOOD
Filer_CF_State: Cluster Enabled  Filer Partner: controller2
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
-----
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
```

**Using the
enable_ontap_pvlin
ks script**

HP-UX 11iv3 does not support ALUA, therefore the default multipathing policy for all the disk storage is active/active. You can use the `enable_ontap_pvlinks` script to change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing the other vendors’ LUNs.

Note_____

Before running `enable_ontap_pvlinks` script, please make sure that Data ONTAP LUNs are visible on the host.

The `enable_ontap_pvlinks` command has the following format:

```
enable_ontap_pvlinks [set | show | unset]
```

Run the following command to set the multipathing policy to active/passive for all Data ONTAP LUNs without disturbing other Vendors’ LUNs:

```
# enable_ontap_pvlinks set
```

Run the following command to display the current multipathing policy for all Data ONTAP LUNs:

```
# enable_ontap_pvlinks show
```

Diagnostic utilities

Diagnostic utilities provided by NetApp

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 compressed 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 customer support.

- ◆ `filer_info`
See “[Collecting filer information](#)” on page 139 for more information.
- ◆ `hpx_info`
See “[Using the HP-UX data-collection tool](#)” on page 140.
- ◆ `switch_info`
switch is the name of the switch you are using (for example, `brocade_info`).
See “[Collecting switch information](#)” on page 141.

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 <code>/opt/NetApp/santools/bin</code> : <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 compressed 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 <code>/opt/NetApp/santools/bin</code>: <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 <code>-d path_and_file_name</code> option enables you to specify a different path in which to place the resulting output. The default path is <code>/tmp/NetApp/ntap_filer_filer_name</code> The <code>-l user:password</code> 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 <code>-n filer_name</code>, where <i>filer_name</i> is the name of the IP address of the storage system (filer) to be interrogated. Result: The output of various status commands are placed in the output directory named <code>ntap_filer_filer_name</code>, where <i>filer_name</i> is replaced by the filer name or IP address. The contents of this directory are placed in a compressed file named <code>ntap_filer_filer_name.tar.gz</code>. 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 compressed file with the following directory path: <pre>/opt/NetApp/santools/ntap_filer_filer1.tar.gz</pre>

Step	Action
3	Send the resulting compressed 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 about 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 compressed file that <code>hpux_info</code> creates. The default location is <code>/tmp/NetApp</code>. <code>-n file_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 compressed 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 compressed 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 `mcddata_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 <code>/opt/NetApp/santools/bin</code> : <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> <i>switch_brand</i> is the name of the switch you are using; for example, “brocade” or “mcdeta”. The <i>-d path_and_file_name</i> option enables you to specify a different path in which to place the resulting output. The default path is <i>//tmp/NetApp/switch_brand_switch_switch_name</i>. The <i>-l user:password</i> option allows the choice of a user name and password on the switch. If this option is not specified, the Brocade default settings of <i>admin:password</i> are used. The target of the command is specified by <i>-n switchname</i>, 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 <i>switch_brand_switch_switchname</i>, 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 compressed file named <i>switch_brand_switch_switchname.tar.gz</i>. Example: This example uses the <i>brocade_info</i> file. The switch <i>brcd1</i> 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 <i>filer1</i> is placed in a compressed file that has the pathname <i>/opt/NetApp/santools/brocade_switch_brcd1.tar.gz</i>.
3	Send the resulting compressed 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 an exceptional long time; 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 the connector has a link and is reporting information such as the correct speed.

What to do: Disconnect the connector and reconnect it.

Index

Symbols

/etc/MANPATH 14

/etc/PATH 14

/opt/NetApp/santools/man 14

A

A/A

Active/Active paths 109

A/PG-C

Active/Passive Group Concurrent paths 109

About

HU 1, 7

UH 1, 7

active paths

Active/Active (A/A) 109

Active/Passive Group Concurrent paths

A/PG-c 109

Array Support Library

see ASL 100

ASL

components 100

definition 100

installing 102

multipathing reports 109

removing 105

Attach Kits

previous kits 2

C

checklist

filer 116

host 117

troubleshooting 116

commands

gunzip 12

igroup create 28

igroup show 29

ioinit (HP-UX) 40

ioscan (HP-UX) 39

lun comment 32

lun setup 29, 30

lun show 37

lvcreate (HP-UX) 46, 51

mount (HP-UX) 48, 53

nwfs (HP-UX) 48, 53

sanlun fcp show adapter 67

sanlun lun show 41, 49, 54, 62

swremove 15

vgcreate (HP-UX) 44

vgdisplay (HP-UX) 46, 51

D

driver

logs off target port 116

E

enable_ontap_pvlinks 3

enter

defined vii

F

FCP HP-UX Host Utilities

contents 2

defined 2

obtaining 12

SAN Toolkit 3, 12

software 12

file system

create 48, 53

mount 48, 53

filer

adding LUNs 26

defined vii

entering commands viii

setup checklist 116

filer_info

diagnostic script 3

G

- gunzip command
 - uncompressing NetApp file 12

H

HBA

- displaying information 67

host

- adding LUNs 26
- create file system 48, 53
- discovering LUNs 38
- displaying HBA information 67
- displaying LUN information 62
- setup checklist 117

HP-UX

- connectivity issues 116
- FCP driver behavior 116
- LVM 38

I

I/O

- numerous retries 116

igroup

- create command 28
- creating 27
- definition 27
- verify 29
- vsa attribute 29

- igroup create command 28

- igroup show command 29

- initiator group, creating 27

- initiator group

- see igroup 27

installation

- NetApp software package 12
- tasks performed 12
- update to search paths 14

- ioinit command 40

- ioscan command 39

L

- logical volumes

- creating 46, 51

- Logical Volume Manager 2

LUN

- defined vii

- lun commands

- lun create 75

- lun map 75

- lun comment command 32

- LUN ID 36

- defined vii

- lun setup command 29

- command line options 30

- interactive script 30

- lun show command 37

LUNs

- adding 26

- creating with command line 30

- creating with interactive script 29

- discovered by host 38

- displaying information on host 62

- lun setup command 30

- tasks for creating on filer 27

- verifying 37

- lvcreate command 46, 51

LVM

- Logcial Volume Manager 38

- used for volume management 2

- volume group entries 43

M

- mount command 48, 53

- multipathing

- LVM 4

- VxVM 4

N

- NetApp software package

- installing 12

- ntap_config_paths utility 45

- reparing paths 120

- nwfs command 48, 53

O

opt/NetApp/santools/bin 14

P

path order

volume groups 45

paths

displaying with sanlun lun show command 62

repairing with ntap_config_paths 120

PVLinks

uses ntap_config_paths 3

S

SAN Toolkit 3

components 3

obtaining 12

san_version

command 3

sanlun fcp show adapter command 67

sanlun lun show command 41, 49, 54, 62

sanlun utility

about 61

locating 26

sanlun fcp show adapter 67

scripts

filer_info 3

ntap_config_paths 3

script used with PVLinks 3

software 12

swremove command

uninstalling software 15

system search path

updated 14

T

troubleshooting

checklists 116

type

defined vii

V

vgcreate command 44

vgdisplay command 46, 51

volume groups

creating logical volumes 46, 51

creating with primary path 44

fixing while offline 118

fixing while online 118

minor numbers 43

path order 44, 45

situations requiring reconfiguration 118

using ntap_conf_paths 45

vsa attribute 29

VxVM

used for volume management 2

W

Where to go for more information 5

