

IBM System Storage N series



HP-UX Host Utilities 6.0 Installation and Setup Guide

Contents

Preface	6
Supported features	6
Websites	6
Getting information, help, and service	6
Before you call	7
Using the documentation	7
Hardware service and support	7
Firmware updates	8
How to send your comments	8
Changes to this document: June 2012	9
The HP-UX Host Utilities	10
Overview of the supported HP-UX environments and protocols	11
How to find instructions for your HP-UX environment	12
Protocols and configurations supported by Host Utilities	12
The FC protocol	13
The FCoE Protocol	13
The iSCSI protocol	13
Supported configurations	14
Features supported by the Host Utilities	15
Multipathing and the HP-UX Host Utilities	15
HBAs and the HP-UX Host Utilities	15
Volume managers and the HP-UX Host Utilities	16
HP-UX configurations that support ALUA	16
SAN booting and the Host Utilities	17
Checklist for planning the Host Utilities installation	18
Installing the HP-UX Host Utilities software	20
Removing the HP-UX Host Utilities software	22
(iSCSI) Installing and Configuring the iSCSI Software Initiator	23
(iSCSI) Installing the iSCSI Software Initiator	23
(iSCSI) Verifying the iSCSI Software Initiator installation	23
(iSCSI) Configuring the iSCSI Software Initiator	24
(iSCSI) Recording the host iSCSI initiator node name	24

(iSCSI) Configuring the discovery targets	25
LUN configuration	27
Overview of LUN configuration and management	27
Discovering LUNs on an HP-UX host	28
(HP-UX 11iv2) Configuring LUNs for use with LVM	31
(HP-UX 11iv3) Configuring LUNs for use with LVM	35
(Veritas) Configuring LUNs for use with VxVM	38
The sanlun utility	43
Displaying host LUN information with sanlun	43
Displaying host HBA information with sanlun	46
How sanlun displays the multipath provider	48
(Veritas) Veritas Dynamic Multipathing configuration	50
(Veritas) What the ASL is	50
(Veritas) ASL array type	50
(Veritas) Information about ASL error messages	51
(Veritas) What the APM is	51
(Veritas) How to get the ASL and APM	52
(Veritas) Installing the ASL and APM software	52
(Veritas) Uninstalling the ASL and APM	54
(Veritas) Upgrading the ASL and APM	54
(Veritas, HP-UX 11iv3) Disabling Native MPIO ALUA	55
(Veritas) The Veritas DMP restore daemon requirements	57
(Veritas) The Veritas DMP I/O policy	57
(Veritas) Displaying multipathing information using sanlun	58
(Veritas) Displaying available paths using VxVM	58
(Veritas, HP-UX 11iV3) Thin provisioning and space reclamation	59
(HP-UX) Setting up SAN boot LUN on HP-UX	61
Troubleshooting	65
Host setup checklist	65
Storage system setup checklist	66
Connectivity problems when using LVM	66
Migrating a configuration from non-ALUA to ALUA without a host reboot	67
(Veritas) Application I/O failures on Veritas VxVM	68
(Veritas) Enclosure-based naming is not working correctly	68
Problems with volume group multipathing	68
Handling a takeover and giveback	69

Events that occur if a host reboots during a takeover	70
Recovering from a host reboot during a takeover	70
Commands you can use to check for problems	72
The HP-UX ioscan command	72
The sanlun utility	72
The enable_ontap_pvlinks script	74
Replacing a LUN on HP-UX 11iv3	74
Copyright information	76
Trademark information	77
Index	80

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the IBM N series support website (accessed and navigated as described in [Websites](#) on page 6).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide IBM N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 6) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Changes to this document: June 2012

This section contains information about the changes made to this guide for HP-UX Host Utilities 6.0. Previously, this guide supported the HP-UX Host Utilities 5.2.

HP-UX Host Utilities 6.0 adds support for more configurations and features; this document has been updated to include information about using those configurations and features.

Any time this document is updated, a note is added to the Release Notes. It is a good practice to check the online Release Notes on a regular basis to determine whether there is new information about using the HP-UX Host Utilities or changes to this guide. The most current versions of the Release Notes and this guide are posted at the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Note: For the most current information about what the HP-UX Host Utilities support, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

June 2012 update

The following information is added to this document:

- The HP-UX Host Utilities support Data ONTAP operating in Cluster-Mode with HP-UX 11iv3. The `sanlun` utility produces output depending on whether the storage systems are running Data ONTAP operating in Cluster-Mode or Data ONTAP operating in 7-Mode. With the HP-UX Host Utilities 6.0 release, the output format of the `sanlun` utility has changed. The format no longer maintains backward compatibility when using LUNs mapped for Data ONTAP operating in 7-Mode. For examples of the output that the `sanlun` command produces, see [The `sanlun` utility](#) on page 43.
- There is a change in process on how you create an igroup. If your controller is running Data ONTAP operating in 7-Mode, you can use the `sanlun fcp show adapter -c` command to get information that you can enter on the controller in order to create an igroup. However, you cannot use this command if you are attempting to create an igroup on a controller running Data ONTAP operating in Cluster-Mode.
- The HP-UX Host Utilities 6.0 installation file name is `ibm_hpunx_host_utilities_6.0_ia_pa.depot.gz`.

The HP-UX Host Utilities

The Host Utilities provide software programs and documentation that you can use to connect your HP-UX host to IBM N series storage systems running Data ONTAP. The software is available as a standard HP-UX depot file.

The Host Utilities include the following components:

- The SAN Toolkit

The toolkit is installed automatically when you install the Host Utilities. This kit provides the following key tools:

Note: This toolkit is common across all configurations and protocols of the HP-UX Host Utilities. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance.

- **(HP-UX 11iv3 only)** The `enable_ontap_pvlinks` tool, which lets you enable the active/passive multipathing policy for Physical Volume Links (PV-Links) with Data ONTAP LUNs.
- **(HP-UX 11iv2 only)** The `ntap_config_paths` tool ensures you have the correct PVlinks multi-pathing configuration to CFO controller LUNs.
- The `sanlun` utility, which helps you to manage Data ONTAP LUNs.
- The `san_version` command, which displays the versions of the Host Utilities.

Note: Previous versions of the Host Utilities also included diagnostics programs. These programs have been replaced by the nSANity Diagnostic and Configuration Data Collector and are no longer installed with the Host Utilities. The nSANity program is not part of the Host Utilities.

See the man pages for these commands for details on using them.

- Documentation

The documentation provides information about installing, setting up, using, and troubleshooting the Host Utilities. The documentation consists of:

- *Installation and Setup Guide*
- *Release Notes*

Note: The *Release Notes* are updated whenever new information about the Host Utilities is available. You should check the *Release Notes* before installing the Host Utilities to see if there is new information about installing and working with the Host Utilities.

- *Quick Command Reference*
- *Host Settings Affected by HP-UX Host Utilities*
- *Quick Start Guide*

Overview of the supported HP-UX environments and protocols

The Host Utilities support several HP-UX environments, such as Native and Veritas.

For details on which environments are supported, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

The following table summarizes key aspects of the main environments:

HP-UX Environment	Notes
Native	<ul style="list-style-type: none"> • Native Multipath I/O (MPIO) <ul style="list-style-type: none"> • This environment uses the HP-UX 11iv3 with native LVM. • Protocols: FC, FCoE, and iSCSI. The FCoE protocol support is for HP-UX 11iv3 March 2011 onwards. • ALUA: Supported with HP UX 11iv3 September 2007 and later • PV-Links <ul style="list-style-type: none"> • HP-UX operating systems: HP-UX 11iv2 with native LVM • Protocols: FC and iSCSI
Veritas	<ul style="list-style-type: none"> • This environment uses Veritas Storage Foundation and its features: <ul style="list-style-type: none"> • Veritas Volume Manager (VxVM) • Veritas Dynamic Multipathing (DMP) • HP-UX operating systems: HP-UX 11iv2 and HP-UX 11iv3 • Required modules: Symantec Array Support Library (ASL) and Array Policy Module (APM) for storage systems • Protocol: FC and FCoE. The FCoE protocol support is for HP-UX 11iv3 March 2011 onwards. • ALUA: Supported with HP UX 11iv3 September 2007 and later • Setup and configuration requirements: <ul style="list-style-type: none"> • You might need to perform some driver setup. • You must install ASL and APM. • If you are using Veritas Storage Foundation with HP-UX 11iv3, you must disable Native MPIO ALUA on Veritas LUNs to ensure that DMP functions properly.

Related information

IBM N series Support page: www.ibm.com/storage/support/nseries/

How to find instructions for your HP-UX environment

Many of the instructions in this guide apply to all the HP-UX environments that the Host Utilities support. In some cases, the commands you use vary based on the environment you are using.

If information or instructions for a step or feature apply only to one environment, this guide notes that fact. To make it easy to quickly identify information for your environment, this guide places a qualifier in the heading to specify that environment. This guide uses the following qualifiers:

Qualifier	The section that follows applies to
(LVM)	Environments using the HP-UX native LVM on either HP-UX 11iv2 or HP-UX 11iv3
(MPIO)	Environments using the HP-UX next-generation mass storage stack that provides native multipathing (MPIO) and agile naming
(VxVM)	Environments using Veritas Volume Manager on either HP-UX 11iv2 or 11iv3
(Veritas)	Environments using Veritas DMP as the multipathing solution
(PV-Links)	Environments using the HP-UX PV-Links on HP-UX 11iv2
(FCoE)	Environments using the Fibre Channel over Ethernet protocol
(FC)	Environments using the Fibre Channel protocol
(iSCSI)	Environments using the iSCSI protocol

There is also information about using the Host Utilities in HP-UX environments in the *Release Notes* and the Host Utilities reference documentation. You can download all the Host Utilities documentation from the [IBM N series Information Center](#).

Protocols and configurations supported by Host Utilities

The Host Utilities provide support for Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI connections to the storage system using direct-attached, fabric-attached, and Ethernet network configurations.

These protocols enable the host to access data on storage systems. The storage systems are targets that have storage target devices called LUNs.

The protocol enables the host to access the LUNs to store and retrieve data.

The sections that follow provide high-level information about these protocols.

Related information

N series interoperability matrix website: www.ibm.com/systems/storage/network/interophome.html

The FC protocol

The FC protocol requires one or more supported HBAs in the host. Each HBA port is an initiator that uses FC to access the LUNs on the storage system. The port is identified by a worldwide port name (WWPN). The storage system uses the WWPNs to identify hosts that are allowed to access LUNs.

You must record the port's WWPN so that you can supply it when you create an initiator group (igroup). You can use the `sanlun fcp show adapter` command to get the WWPN.

When you create the LUN, you must map it to that igroup. The igroup then enables the host to access the LUNs on the storage system using the FC protocol based on the WWPN.

For more information about using FC with your storage system, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP.

Related tasks

Displaying host HBA information with `sanlun` on page 46

Related references

Overview of LUN configuration and management on page 27

The FCoE Protocol

Fibre Channel over Ethernet (FCoE) is a new model for connecting hosts to storage systems. Like the traditional FC protocol, FCoE maintains existing FC management and controls, but it uses a 10-gigabit Ethernet network as the hardware transport.

Setting up an FCoE connection requires one or more supported Converged Network Adapters (CNAs) in the host, connected to a supported Data Center Bridging (DCB) Ethernet switch. The CNA is a consolidation point and effectively serves as both an HBA and an Ethernet adapter.

The iSCSI protocol

The iSCSI protocol is implemented on both the host and the storage system.

On the host, the iSCSI protocol is implemented over the host's standard gigabit Ethernet interfaces using a software driver.

The HP-UX host does not support hardware iSCSI HBAs.

On the storage system, the iSCSI protocol can be implemented over the storage system's standard Ethernet interface using a software driver that is integrated into Data ONTAP.

The connection between the initiator and target uses a standard TCP/IP network. The network can be a dedicated TCP/IP network, or it can be your regular public network; however, it is best to use a private network to transfer data between the host and the storage system. The storage system listens for iSCSI connections on IP port 3260.

Each host has a single iSCSI node name for all iSCSI ports. You need to make a note of the iSCSI node name so that you can supply it when you create an igroup. The storage system identifies hosts that are allowed to access LUNs based on the iSCSI initiator node name that you supplied when you created the igroup.

For more information about using iSCSI with your storage system, see the *SAN Administration Guide (called Data ONTAP Block Access Management Guide for iSCSI and FC in Data ONTAP 8.1 and earlier)* for your version of Data ONTAP.

Related tasks

[\(iSCSI\) Recording the host iSCSI initiator node name](#) on page 24

Related references

[Overview of LUN configuration and management](#) on page 27

Supported configurations

The Host Utilities support fabric-attached, direct-attached, and network-attached configurations.

The Host Utilities support the following basic configurations:

- Fabric-attached storage area network (SAN)/Fibre Channel over Ethernet network. The Host Utilities support two variations of fabric-attached SANs:
 - A single-host FC connection from the HBA to the storage system through a single switch. A host is cabled to a single FC switch that is connected by cable to redundant FC ports on a high availability storage system configuration. A fabric-attached, single-path host has one HBA.
 - Two or more FC connections from the HBA to the storage system through dual switches or a zoned switch. In this configuration, the host has at least one dual-port HBA or two single-port HBAs. The redundant configuration avoids the single point of failure of a single-switch configuration. This configuration requires that multipathing be enabled.

Note: Use redundant configurations with two FC switches for high availability in production environments. However, direct FC connections and switched configurations using a single, zoned switch might be appropriate for less critical business applications.

- FC direct-attached. A single host with a direct FC connection from the HBA to stand-alone or active/active storage system configurations.
- iSCSI network-attached. In an iSCSI environment, all methods of connecting Ethernet switches to a network that have been approved by the switch vendor are supported. Ethernet switch counts is not a limitation in Ethernet iSCSI topologies. Refer to the Ethernet switch vendor documentation for specific recommendations and best practices.

The *SAN Configuration Guide (called Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier)* provides detailed information, including diagrams, about the supported

topologies. There is also configuration information in the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP. Refer to those documents for complete information about configurations and topologies.

Features supported by the Host Utilities

Some of the supported features include the following:

- Multiple paths to the storage system when a multipathing solution is installed (PV-Links, DMP, Native MPIO)
- HBAs
- Volume managers (VxVM, LVM)
- ALUA
- SAN booting

The N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6) describes the features supported for respective configurations.

Related information

N series interoperability matrix website: www.ibm.com/systems/storage/network/interphome.html

Multipathing and the HP-UX Host Utilities

The HP-UX Host Utilities support different multipathing solutions based on your configuration.

Using multipathing allows you to configure multiple network paths between the host and storage system. That way, if one path fails, traffic continues on the remaining paths. For a host to have multiple paths to a LUN, you must have multipathing enabled.

The HP-UX Host Utilities support PV-Links, DMP, and Native MPIO multipathing solutions.

LVM uses PV-Links and native MPIO to provide alternative paths if a problem causes the active path to disappear.

VxVM uses DMP to provide multipathing. If you want to use VxVM to manage your LUNs, you must install the Symantec ASL and APM with Veritas Storage Foundation for storage systems.

HBAs and the HP-UX Host Utilities

The HP-UX Host Utilities support a number of HBAs.

Ensure the supported HBAs are installed before you install the Host Utilities.

Note: For details on the specific HBAs that are supported and the required firmware and FC drivers, see the IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

Related information

IBM N series interoperability matrix website: www.ibm.com/systems/storage/network/interphome.html

Volume managers and the HP-UX Host Utilities

The HP-UX Host Utilities support different volume management solutions based on your environment.

The HP-UX Host Utilities manage the volumes using either HP-UX LVM or Veritas VxVM. In some cases, you might set up your host to use both LVM and VxVM.

HP-UX configurations that support ALUA

The HP-UX Host Utilities support ALUA in environments using the FC protocol with Native MPIO as long as both your version of the HP-UX operating system and Data ONTAP support ALUA. Certain environments running Veritas Storage Foundation also support ALUA.

ALUA defines a standard set of SCSI commands for discovering and managing multiple paths to LUNs on FC and iSCSI SANs. You should enable ALUA when your Host Utilities configuration supports it. ALUA is enabled on the igroup mapped to IBM N series LUNs that are used by the HP-UX host.

The following table provides information about which versions of HP-UX using Native MPIO and which versions of Data ONTAP support ALUA:

HP-UX version	ALUA support	Minimum Data ONTAP version for ALUA
HP UX 11iv3 September 2007 and later	Yes Note: ALUA is mandatory with this version of HP-UX.	7.2.5 or later
HP-UX 11iv3 February 2007 release	No	Not applicable
HP-UX 11iv2	No	Not applicable

If you are using Veritas Storage Foundation 5.0.1 with HP-UX 11iv3, you must disable Native MPIO ALUA on Veritas LUNs to ensure that DMP functions properly. Otherwise, the `sanlun` utility does not correctly display information about the DMP node. For information about disabling ALUA, see the Symantec TechNote—*How to Disable HP-UX 11iv3 Native Multi-Pathing ALUA mode for Storage Foundation 5.0 and 5.0.1*.

For information about which combinations of HP-UX, Data ONTAP, and Veritas Storage Foundation are supported with which versions of the Host Utilities, see the IBM N series

interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

Related tasks

[\(HP-UX 11iv3\) Configuring LUNs for use with LVM](#) on page 35

[Discovering LUNs on an HP-UX host](#) on page 28

[Migrating a configuration from non-ALUA to ALUA without a host reboot](#) on page 67

Related references

[The enable_ontap_pvlinks script](#) on page 74

Related information

[IBM N series interoperability matrix website: *www.ibm.com/systems/storage/network/interophome.html*](#)

[How to Disable HP-UX 11iv3 Native Multi-Pathing ALUA mode for Storage Foundation 5.0 and 5.0.1](#)

SAN booting and the Host Utilities

The Host Utilities support SAN booting with FC and FCoE protocols in HP-UX environments. SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an HP-UX host.

Configuring SAN booting on a storage system LUN allows you to:

- Remove the hard drives from your servers and use the SAN for booting needs, eliminating the costs associated with maintaining and servicing hard drives
- Consolidate and centralize storage
- Use the reliability and backup features of the storage system

The downside of SAN booting is that loss of connectivity between the host and storage system can prevent the host from booting. Be sure to use a reliable connection to the storage system.

Checklist for planning the Host Utilities installation

Installing the Host Utilities and setting up your system involves numerous tasks that are performed on both the storage system and the host. The checklist provides a high-level overview of these tasks.

If you are an experienced HP-UX user, this checklist can serve as a quick start guide to installing and setting up the Host Utilities.

The detailed steps for each of the tasks presented in the checklist are provided in the *HP-UX Host Utilities Quick Start Guide*.

Note: Occasionally, there are known problems that can affect your system setup. Read the *Host Utilities Release Notes* before you install the Host Utilities. The *Release Notes* are updated whenever an issue is found and might contain information about the Host Utilities that was observed after this guide was produced.

Task 1: Ensure the prerequisites for installing and setting up the Host Utilities have been met

1. Verify that your system setup is correct.

Check the IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6) for the most current information about system requirements.

2. Verify that your storage system has Data ONTAP installed and is running correctly with the licensed protocol for your environment.

3. (FC/FCoE) If you are using a switch, verify that it is:

- Set up correctly.
- Zoned appropriately.
- Cabled correctly according to the instructions in the *SAN Configuration Guide (called Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier)* for your version of Data ONTAP.
- Powered on in the correct order: switch, disk shelves, storage systems, and then the host.

Note: For information about supported topologies for your version of Data ONTAP, see *SAN Administration Guide (called Block Access Management Guide for iSCSI and FC in Data ONTAP 8.1 and earlier)*.

4. Confirm that the host and the storage system can communicate by verifying the following:

- Host ports have "logged in" on controller.
- LUNs are visible if mapped to the host.

Task 2: Install the Host Utilities

1. Download a copy of the compressed file containing the Host Utilities from the IBM N series support website (accessed and navigated as described in [Websites](#) on page 6).

2. Uncompress the file and extract the SAN Toolkit software package.
You can use `gunzip` command to uncompress the file.

```
zcat ibm_aix_host_utilities_6.0.tar.Z | tar -xvf -
```

3. Install the Host Utilities software by using the `swinstall` command.

Note: If you are using iSCSI, make sure you install the iSCSI Software Initiator before you install the Host Utilities.

4. Log out and log back in to enable the updates of the installation script.

(iSCSI) Task 3: Configure the iSCSI protocol

1. Download the iSCSI software initiator from the [HP ITRC website](#).

2. Use the `swlist` command to verify that the file is downloaded correctly and then install the initiator using the `swinstall` file.

3. Configure the iSCSI Software Initiator by using `iscsiutil` command.

For more information, see the instructions provided with the iSCSI Software Initiator.

Note: If you are using multipathing, you must follow the instructions provided by the multipathing vendor to set it up to work with iSCSI.

4. Record the iSCSI node name of the host.

5. Set up target discovery.

6. **Optional:** Set up Challenge Handshake Authentication Protocol (CHAP) on the host and the storage system by following the instructions provided in the HP-UX documentation.

Task 4: Set up access between the host and the LUNs on the storage system

1. Create and map igroups and LUNs.

If your environment supports ALUA, make sure ALUA is enabled.

2. Discover the new LUNs.

You can use the `ioscan` command to discover the LUNs and the `insf` command to create device special files for the newly discovered LUNs.

3. Configure LUNs to work with the volume manager.

4. Display information about the LUNs.

You can use the `sanlun` command to display information about the LUNs and the HBAs.

Installing the HP-UX Host Utilities software

Before you begin

- Host system meets the Host Utilities requirements.
- If you are upgrading the HP-UX Host Utilities 6.0 from an earlier version, you should have uninstalled the earlier version.
- **(iSCSI only)** Installed your iSCSI software initiator before you install the Host Utilities.

Steps

1. Download the HP-UX Host Utilities file `ibm_hpux_host_utilities_6.0_ia_pa.depot.gz` to your HP-UX host.
2. Uncompress the `ibm_hpux_host_utilities_6.0_ia_pa.depot.gz` file by entering the following command:

```
# gunzip ibm_hpux_host_utilities_6.0_ia_pa.depot.gz
```

The system places the extracted software in the directory to which you uncompressed the depot file.

3. Install the software by entering the following command:

```
swinstall -s /depot_path
```

`depot_path` provides the path and name of the depot file.

The `swinstall` command runs an installation script that verifies the status of your HP-UX setup. If your system meets the requirements, this script installs the `sanlun` utility and diagnostic scripts in the `/opt/Ontap/santools/bin` directory.

Example

The following output is similar to what you might see when you install the software for the first time. This example assumes that you uncompressed the depot file to the `/` directory on an HP-UX host called `hpux_95`.

```
===== 10/20/11 14:35:16 IST BEGIN swinstall SESSION
(non-interactive) (jobid=hpux_24-0044)

* Session started for user "root@hpux_24".

* Beginning Selection
* Target connection succeeded for "hpux_24:/".
* Source:
  /ibm_hpux_host_utilities_5.2_ia_pa.depot
* Targets:
  hpux_24:/
* Software selections:
  Ontap_santoolkit.command_itanium,r=5.2,fr=5.2,fa=HP-UX_B.11.22_IA
  Ontap_santoolkit.command_parisc,r=5.2,fr=5.2,fa=HP-UX_B.11.00_32/64
  Ontap_santoolkit.man,r=5.2
  Ontap_santoolkit.support_scripts,r=5.2
* Selection succeeded.

* Beginning Analysis and Execution
```

```
* Session selections have been saved in the file
  "/.sw/sessions/swinstall.last".
* The analysis phase succeeded for "hpux_24:/".
* The execution phase succeeded for "hpux_24:/".
* Analysis and Execution succeeded.

NOTE:   More information may be found in the agent logfile using the
        command "swjob -a log hpux_24-0044 @ hpux_24:/".

===== 10/20/11 14:35:38 IST  END swinstall SESSION (non-interactive)
         (jobid=hpux_24-0044)
```

Related tasks

[\(iSCSI\) Installing the iSCSI Software Initiator](#) on page 23

[Removing the HP-UX Host Utilities software](#) on page 22

Removing the HP-UX Host Utilities software

You can uninstall the Host Utilities software by using the `swremove` command.

About this task

When upgrading to HP-UX Host Utilities 6.0 from versions earlier than 5.1, you should first uninstall the earlier version. The reason is that starting with HP-UX Host utilities 5.1, diagnostic utilities that were previously packaged with Host Utilities software is removed. The only way to remove those utilities is to uninstall the current Host Utilities software.

Step

1. Remove the Host Utilities software by using the `swremove` command.

```
# swremove Ontap_santoolkit
```

Example (Host Utilities)

The following example displays the output you get when you execute the `swremove` command to remove the Host Utilities software from an HP-UX host called `hpux_95`:

```
===== 10/20/11 14:35:46 IST BEGIN swremove SESSION
(non-interactive) (jobid=hpux_24-0045)

* Session started for user "root@hpux_24".

* Beginning Selection
* Target connection succeeded for "hpux_24:/".
* Software selections:
  Ontap_santoolkit.command_itanium,l=/opt/Ontap/santools,r=5.2,fr=5.2,fa=HP-UX_B.11.22_IA
  Ontap_santoolkit.man,l=/opt/Ontap/santools,r=5.2
  Ontap_santoolkit.support_scripts,l=/opt/Ontap/santools,r=5.2
* Selection succeeded.

* Beginning Analysis
* Session selections have been saved in the file
"/.sw/sessions/swremove.last".
* The analysis phase succeeded for "hpux_24:/".
* Analysis succeeded.

* Beginning Execution
* The execution phase succeeded for "hpux_24:/".
* Execution succeeded.

NOTE: More information may be found in the agent logfile using the
command "swjob -a log hpux_24-0045 @ hpux_24:/".

===== 10/20/11 14:35:48 IST END swremove SESSION (non-interactive)
(jobid=hpux_24-0045)
)
```

(iSCSI) Installing and Configuring the iSCSI Software Initiator

When you are using the iSCSI protocol, you must perform some additional tasks to complete the installation of the Host Utilities.

(iSCSI) Installing the iSCSI Software Initiator

After you download the iSCSI Software initiator, you must verify that the file is downloaded correctly and then run the `swinstall` command on the HP-UX host to install the iSCSI Software Initiator.

Steps

1. Log in as root.
2. Verify that the file is downloaded correctly by entering the following command:

```
# swlist @ /depot_path
```

3. Install the iSCSI Software Initiator by entering the following command:

```
# swinstall -x autoreboot=true -s /depot_path
```

`depot_path` provides the path and name of the depot file.

The `autoreboot=true` option causes a system reboot when the installation is complete.

For more information, see the following:

- The *HP-UX iSCSI Software Initiator Release Notes* provide information about patch dependencies.
- The *HP-UX iSCSI Software Initiator Support Guide* provides instructions for configuring the iSCSI subsystem.

(iSCSI) Verifying the iSCSI Software Initiator installation

To verify the iSCSI Software Initiator installation, you can use the `swlist` or `ioscan` commands.

Step

1. Verify the iSCSI Software Initiator installation by entering the following command:

```
# swlist iSCSI-00
```

```
iSCSI-00 B.11.31.01 HP-UX iSCSI Software Initiator
iSCSI-00.iSCSI-SWD B.11.31.01 HP-UX iSCSI Software Initiator
```

Alternatively, you can use the `ioscan -fnC iscsi` command to verify the iSCSI Software Initiator's installation.

```
Class I H/W Path Driver S/W State H/W Type Description
=====
iscsi 1 255/0 iscsi CLAIMED VIRTBUS iSCSI Virtual Node
```

(iSCSI) Configuring the iSCSI Software Initiator

To configure the iSCSI Software Initiator, you must add the path of several executables to the root path and run the `iscsiutil` command.

Steps

1. Add the path of `iscsiutil` and other `iscsi` executables to the root path:

```
PATH=$PATH:/opt/iscsi/bin
```

2. Configure iSCSI Software Initiator name by entering the following command:

```
# iscsiutil [iscsi-device-file] -i -N iscsi-initiator-name
```

iscsi-device-file is the iSCSI device file path `/dev/iscsi`. This argument is optional when other options, such as `-i` and `-N`, are included in the command.

- `-i` configures iSCSI Software Initiator information.
- `-N` is the initiator name option. When preceded by `-i`, it requires the iSCSI Software Initiator name as an argument. The first 256 characters of the name string are stored in the iSCSI persistent information.
- *iscsi-initiator-name* is the Software Initiator name you have chosen, in the `iqn` format.

(iSCSI) Recording the host iSCSI initiator node name

You need to supply the iSCSI initiator node name when you create `igroups` on the storage system. It is a good practice to record the node name before you create the `igroups`. You can use the `iscsiutil -l` command to display the node name.

Steps

1. Display the iSCSI node name by entering the following command on the host:

```
# iscsiutil -l
```

Example

This command displays information about your iSCSI setup. The line containing the node name is shown in **bold**.

```
# iscsiutil -l
Initiator Name           : iqn.1986-03.com.hp:hpux11
Initiator Alias          :
Authentication Method    :
CHAP Method              : CHAP_UNI
Initiator CHAP Name      :
CHAP Secret              :
NAS Hostname             :
NAS Secret               :
Radius Server Hostname   :
Header Digest            : None,CRC32C (default)
Data Digest              : None,CRC32C
```

This output has been truncated to make the document easier to read.

- Record the node name so that you will have it when you create igroups.

(iSCSI) Configuring the discovery targets

You need to configure the iSCSI initiator to discover the target so that the host can access LUNs on the target.

Steps

- Add one or more discovery targets by entering the following command:

```
iscsiutil [/dev/iscsi] -a -I ip-address [-P tcp-port] [-M portal-grp-tag]
```

- a* adds a discovery target address into iSCSI persistent information. Only discovery target addresses can be added using this option.
- ip-address* is the IP address or host name component of the target network portal.
- portal-grp-tag* is the target portal group tag. The default target portal group tag for discovery targets is 1.
- tcp-port* is the TCP port component of the discovery target network portal. The default iSCSI TCP port number is 3260.

- View the configured discovery targets by entering the following command:

```
# iscsiutil -p -D
```

- Discover the operational target devices by entering the following command:

```
# /usr/sbin/ioscan -H 255/0
```

- Create the device file for the targets by entering the following command:

```
# /usr/sbin/insf -H 255/0
```

5. Display the operational targets by entering the following command:

```
# iscsiutil -p -O
```

LUN configuration

Configuring LUNs involves several tasks. Whether you are executing the Host Utilities in an HP-UX LVM or VxVM environment determines the tasks you must perform.

Overview of LUN configuration and management

LUN configuration and management in an HP-UX environment involves several tasks. Some tasks do not apply to all environments.

Task	Description
Zoning	Ensure that you have zoned the HP-UX host and the controllers correctly.
Creating and mapping igroups and LUNs	<p>An igroup allows the storage system to identify which hosts can access LUNs on that storage system. To create an igroup you need either of the following:</p> <ul style="list-style-type: none"> • (FC/FCoE) The WWPN for the HBA ports on the host. You can use the <code>sanlun fcp show adapter</code> command to get the WWPN. • (iSCSI) The iSCSI initiator node name for the host. You can use the <code>iscsiutil -l</code> command to get this. <p>After you create the igroup, you must create LUNs on the storage system and map the LUNs to the igroup.</p> <p>The <i>SAN Configuration Guide (Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier)</i> for your version of Data ONTAP provides information about creating igroups and LUNs.</p>
Enabling ALUA	If your environment supports ALUA, make sure ALUA is enabled. For more information about enabling ALUA, see <i>Data ONTAP Commands: Manual Page Reference for Cluster-Mode</i> for Cluster-Mode and <i>Data ONTAP Commands: Manual Page Reference for 7-Mode, Volume 2</i>

Task	Description
Discovering LUNs	After you create the LUN and map it to your igroup, you must discover the LUN as a host device.
Configuring volume management software	You must configure the LUNs so they are under the control of a volume manager (LVM or VxVM) that is supported by your Host Utilities environment.

Related tasks

[Displaying host HBA information with `sanlun`](#) on page 46

[\(iSCSI\) Recording the host iSCSI initiator node name](#) on page 24

Discovering LUNs on an HP-UX host

To configure and manage the LUNs, the LUNs must be discovered by the host. Rebooting the host automatically discovers new LUNs. You have to perform a set of steps if rebooting is not a reasonable action.

Steps

1. Log in as root on the host.
2. To discover the newly created LUNs, perform an `ioscan` on the HP-UX host by entering the following command:

```
# ioscan -fnC disk
```

The output from the `ioscan` command provides one of the following information:

- No device special files for LUNs exist 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.

Example

The output from the `ioscan` command provides one of the following information

Following is the output of command, using the FC protocol:

```
# ioscan -fnC disk
Class   I  H/W Path      Driver S/W State  H/W Type  Description
-----
disk    2  0/4/0/0/0/0.58.11.0.0.0.0  sdisk  CLAIMED    DEVICE    ONTAP  LUN
        /dev/dsk/c5t0d0 /dev/rdisk/c5t0d0
disk    0  0/4/0/0/0/0.58.15.0.0.0.0  sdisk  CLAIMED    DEVICE    ONTAP  LUN
        /dev/dsk/c1t0d0 /dev/rdisk/c1t0d0
disk    3  0/4/0/0/0/1.59.11.0.0.0.0  sdisk  CLAIMED    DEVICE    ONTAP  LUN
        /dev/dsk/c7t0d0 /dev/rdisk/c7t0d0
```

```
disk      1  0/4/0/0/0/1.59.15.0.0.0.0  sdisk  CLAIMED  DEVICE  ONTAP  LUN
          /dev/dsk/c3t0d0           /dev/rdisk/c3t0d0
```

This example illustrates the output that is displayed in an iSCSI environment:

```
# ioscan -fnC disk

disk 6 255/0/0.0.0.1 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d1 /dev/rdisk/c2t0d1
disk 7 255/0/0.0.0.2 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d2 /dev/rdisk/c2t0d2
disk 8 255/0/0.0.0.3 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d3 /dev/rdisk/c2t0d3
disk 9 255/0/0.0.0.4 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d4 /dev/rdisk/c2t0d4
disk 10 255/0/0.0.0.5 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d5 /dev/rdisk/c2t0d5
disk 11 255/0/0.0.0.6 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d6 /dev/rdisk/c2t0d6
disk 12 255/0/0.0.0.7 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t0d7 /dev/rdisk/c2t0d7
disk 13 255/0/0.0.1.0 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t1d0 /dev/rdisk/c2t1d0
disk 14 255/0/0.0.1.1 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t1d1 /dev/rdisk/c2t1d1
disk 15 255/0/0.0.1.2 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t1d2 /dev/rdisk/c2t1d2
disk 16 255/0/0.0.1.3 sdisk CLAIMED DEVICE ONTAP LUN
      /dev/dsk/c2t1d3 /dev/rdisk/c2t1d3
```

3. Install special files and create device special files on the host by entering the following command:

```
# ioinit -i
```

If the `ioinit -i` command does not create device special files, install the device special files by entering the `insf -e` command.

Example

This example uses the `ioinit -i` to install the special files.

```
# 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
```

4. Verify the information about the device special files by entering the following command:

```
# sanlun lun show -p all
```

Example

The `sanlun lun show -p all` command displays the following type of information if you are using Data ONTAP operating in 7-Mode:

```
# sanlun lun show -p all

ONTAP Path: f3070-210-38:/vol/vol1/lun1
```

```

                LUN: 0
                LUN Size: 2g
                Controller CF State: Cluster Enabled
                Controller Partner: f3070-210-37
                Host Device: /dev/rdisk/disk5
                Mode: 7
                Multipath Provider: None
-----
host    controller /dev/dsk
path    path       filename
state   type       or hardware path
-----
up      primary    /dev/dsk/c1t0d0    fcd2    1a
up      primary    /dev/dsk/c3t0d0    fcd3    1b
up      secondary  /dev/dsk/c5t0d0    fcd2    1b
up      secondary  /dev/dsk/c7t0d0    fcd3    1a

```

Example

The `sanlun lun show -p all` command displays the following type of information if you are using Data ONTAP operating in Cluster-Mode:

```

# sanlun lun show -p all
                ONTAP Path: vs39:/vol/vol24_3_0/lun24_0
                LUN: 12
                LUN Size: 3g
                Host Device: /dev/rdisk/disk208
                Mode: C
                Multipath Provider: None
-----
host    vserver    /dev/dsk
path    path       filename
state   type       or hardware path
-----
up      primary    /dev/dsk/c3t1d4    fcd2    fc1_0    0
up      primary    /dev/dsk/c16t1d4   fcd2    fc1_1    0
up      primary    /dev/dsk/c130t1d4  fcd5    fc2_0    0
up      primary    /dev/dsk/c135t1d4  fcd5    fc2_1    0
up      secondary  /dev/dsk/c8t1d4    fcd2    fc6_0    1
up      secondary  /dev/dsk/c11t1d4   fcd2    fc6_1    1
up      secondary  /dev/dsk/c13t1d4   fcd2    fc8_0    1
up      secondary  /dev/dsk/c25t1d4   fcd2    fc8_1    1

```

Note: If you are using HP-UX 11iv2, you do not need to perform the rest of the steps.

- (HP-UX 11iv3 February 2007 release)** If necessary, use the `enable_ontap_pvlinks` script to change the multipathing policy to an active/passive configuration for Data ONTAP LUNs without disturbing LUNs from other vendors by entering the command:

```
# enable_ontap_pvlinks set
```

Note: The HP-UX 11iv3 February 2007 release does not support ALUA, so the default multipathing policy for all the disk storage is an active/active configuration. If you are using the HP-UX 11iv3 February 2007 release, you might need to change the multipathing policy to an active/passive configuration for Data ONTAP LUNs.

- (HP-UX 11iv3 February 2007 release)** You can verify the current multipathing policy for all Data ONTAP LUNs by entering the following command:

```
# enable_ontap_pvlinks show
```

(HP-UX 11iv2) Configuring LUNs for use with LVM

You must perform several tasks to set up LUNs for an LVM environment.

Steps

1. Create a physical volume on the LUN by entering the following command:

```
# pvcreate /dev/rdisk/path_disk_device
```

path_disk_device is the path to the disk device that contains the LUN.

Example

The `pvcreate` command produces output similar to the following example:

```
# pvcreate /dev/rdisk/c10t0d4
Physical volume "/dev/rdisk/c10t0d4" has been successfully created.
```

You must perform this process at least one time for each LUN, when you have multiple paths to that LUN.

2. Check the minor numbers of all existing volume groups by entering the following command:

```
# ls -l /dev/*/group
```

LVM requires a unique minor number for each volume group device entry.

Example

The `ls -l /dev/*/group` command provides output similar to the example, where the next unused minor device number is 0x010000.

```
# ls -la /dev/*/group
crw-r----- 1 root sys 64 0x000000 Dec 11 19:24 /dev/vg00/group
```

3. Create a `/dev` entry by entering the following command:

```
# mkdir /dev/directory_name
```

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.

```
# mkdir /dev/vg_ntap01
```

4. Create a device node on the host for the physical volume group. You must supply a unique minor device number. Do not enter a number that another volume group is using.

Example

This example creates a device node using the unique minor number 0x010000.

```
# mknod /dev/vg_ntap01/group c 64 0x010000
```

Note: The device node you create is used only for volume group and logical volume operations; it is not used for I/O.

5. Create a volume group with one primary path by entering the following command:

```
# vgcreate vg_name /dev/dsk/path_disk_device
```

vg_name is a volume group name, with or without the */dev/* prefix.

path_disk_device is the path name to the disk.

Note: Use a regular or block disk device node. Do not use raw disk devices.

Example

In this example, the `vgcreate` command sets up a primary path to the volume group.

```
# 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
```

6. Add the remaining paths to the volume group using the `ontap_config_paths` utility, which is in the `/opt/Ontap/santools/bin/` directory.

`ontap_config_paths`

Example

The `ontap_config_paths` utility adds the rest of the paths using the `vgextend` command.

```
# ontap_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
#
```

You can verify the path by entering the following command:

```
# vdisplay -v volume_group
```

Note: The `ontap_config_paths` utility adds missing paths to the LVM's volume group and places the paths in the correct order. It specifies the disk device special files that correspond to the primary paths first; then it specifies the device special files that correspond to the secondary paths.

7. Create one or more logical volumes in the volume group by entering the following command:

```
# lvcreate [-L LogicalVolumeSize | -l LogicalExtentsNumber] volume_group
```

LogicalVolumeSize with an uppercase `-L` option is the total size of the logical volume in megabytes.

LogicalExtentsNumber with a lowercase `-l` option is the number of physical extents that you want to assign to the logical volume.

The value you supply for size must not equal or exceed the available free physical extents in the volume group. The LUN reserves an area for physical volume, volume group, and LVM information.

The `lvcreate` and the `lvextend` commands fail if the size you specify here does not fit into the volume group.

You can check the available free physical extents by entering the following command:

```
# vgdisplay -v volume_group
```

Note: After you create a logical volume, PV-Links path management is in effect on the logical volume.

Example

The following example creates a logical volume for `dev/vg_ntap01/` that uses all the physical extents available for that volume. The `vgdisplay` command displays the number of available physical extents (511). The `lvcreate` command with a lowercase `-l` option assigns them to the logical volume. The man pages provide more information about `vgdisplay` and `lvcreate`.

```
# vgdisplay 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/lvol11" has been successfully created with
character device "/dev/vg_ntap01/rlvol11".
Logical volume "/dev/vg_ntap01/lvol11" has been successfully extended.
```

```
Volume Group configuration for /dev/vg_ntap01 has been saved in
/etc/lvmconf/vg_ntap01.conf
#
```

8. Create a file system on the logical volume device special file by entering the following command:

```
# newfs -F vxfs /dev/volume_group/character_device
character_device is the character device listed in the output of the vdisplay -v
volume_group command.
```

Example

The following example creates a file system on the logical volume /dev/vg_ntap01:

```
# 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
```

9. After you create the new file system, create a mount point by entering the following command:

```
# mkdir /mnt/vg_ntap01_lv1
```

10. Mount the new file system by entering the following command:

```
# mount logical_volume mount_point
```

Example

This example mounts the file system you created in the previous step.

```
# 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
#
```

11. Verify the LUN path priority by entering the following command:

```
# sanlun lun show -p all
```

Example

This example shows the type of output you see using the sanlun command.

```
# sanlun lun show -p
ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
LUN: 1
LUN Size: 100m
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-38
Mode: 7
VG: /dev/vg_ntap01
Multipath Policy: A/P
Multipath Provider: Native
-----
host      controller /dev/dsk          controller      PVlinks
path      path         filename          target          path failover
state     type         or hardware path  adapter        port           priority
-----
up        primary     /dev/dsk/c9t0d1   fcd1           1a             0
up        primary     /dev/dsk/c4t0d1   fcd0           1b             1
```

up	secondary	/dev/dsk/c11t0d1	fcd1	1b	2
up	secondary	/dev/dsk/c6t0d1	fcd0	1a	3

(HP-UX 11iv3) Configuring LUNs for use with LVM

Before you begin

Before mapping the LUNs, enable ALUA on the igroup.

Steps

1. Create a physical volume on the LUN by entering the following command:

```
# pvcreate /dev/rdisk/path_disk_device
path_disk_device is the path to the disk device that contains the LUN.
```

Example

The `pvcreate` command displays an output similar to the following example:

```
# pvcreate /dev/rdisk/disk56
Physical volume "/dev/rdisk/disk56" has been successfully created.
```

2. Create a `/dev` entry by entering the following command:

```
# mkdir /dev/directory_name
```

Example

```
# mkdir /dev/vg_ntap01
```

3. Create a volume group by entering the following command:

```
# vgcreate vg_name /dev/disk/path_disk_device
```

vg_name is a volume group name, with or without the `/dev/` prefix.

path_disk_device is the path name to the disk.

Note: Use a regular or block disk device node. Do not use raw disk devices.

Example

In this example, the `vgcreate` command creates a volume group by name `vg_ntap01`.

```
# vgcreate /dev/vg_ntap01 /dev/disk/disk56
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
```

4. Verify the path by entering the following command:

```
# vgsdisplay -v volume_group
```

5. Create one or more logical volumes in the volume group by entering the following command:

```
# lvcreate [-L LogicalVolumeSize | -l LogicalExtentsNumber] volume_group
```

LogicalVolumeSize with an uppercase `-L` option is the total size of the logical volume in megabytes.

LogicalExtentsNumber with a lowercase `-l` option is the number of physical extents that you want to assign to the logical volume. The value you supply for size must not equal or exceed the available free physical extents in the volume group.

The LUN reserves an area for physical volume, volume group, and LVM information. The `lvcreate` and the `lvextend` commands fail if the size you specify here does not fit into the volume group.

You can check the available free physical extents by entering the following command:

```
# vdisplay -v volume_group
```

Example

The following example creates a logical volume for `dev/vg_ntap01/` that uses all the physical extents available for that volume. The `vdisplay` command displays the number of available physical extents (511). The `lvcreate` command with a lowercase `-l` option assigns them to the logical volume. The man pages provide more information about `vdisplay` and `lvcreate`.

```
# vdisplay vg_ntap01
--- Volume groups ---
VG Name                /dev/vg_ntap01
VG Write Access        read/write
VG Status              available
Max LV                 255
Cur LV                2
Open LV               2
Max PV                 255
Cur PV                20
Act PV                20
Max PE per PV         1016
VGDA                   40
PE Size (Mbytes)      4
Total PE              10220
Alloc PE               6000
Free PE               4220
Total PVG              0
Total Spare PVs        0
Total Spare PVs in use 0
VG Version             2.0
VG Max Size            1t
VG Max Extents         4096

# 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
#
```

6. Create a file system on the logical volume device special file by entering the following command:

```
# newfs -F vxfs /dev/volume_group/character_device
character_device is the character device listed in the output of the vgdisplay -v
volume_group command.
```

Example

The following example shows how to create a file system on the logical volume /dev/vg_ntap01/rlvol1:

```
# 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
#
```

7. After you create the new file system, create a mount point by entering the following command:

```
# mkdir /mnt/vg_ntap01_lv1
```

8. Mount the new file system by entering the following command:

```
# mount logical_volume mount_point
```

Example

This example is the output of the file system that you mounted in the previous step.

```
# 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
#
```

9. Verify the LUN path priority by entering the following command:

```
# sanlun lun show -p
```

This example shows the type of output you see using the `sanlun` command for Data ONTAP operating in 7-Mode.

```
# sanlun lun show -p all
ONTAP Path: f3070-210-38:/vol/vol1/lun1
LUN: 0
LUN Size: 2g
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-37
Host Device: /dev/rdisk/disk5
Mode: 7
VG: /dev/vg01
Multipath Policy: A/A
Multipath Provider: Native
-----
host controller /dev/dsk          controller HP A/A
path path filename             host target path failover
state type or hardware path adapter port priority
-----
up primary /dev/dsk/c1t0d0             fcd2 1a 0
```

```

up      primary   /dev/dsk/c3t0d0    fcd3  1b      0
up      secondary /dev/dsk/c5t0d0    fcd2  1b      1
up      secondary /dev/dsk/c7t0d0    fcd3  1a      1

```

This example shows the type for output you see using the `sanlun` command for Data ONTAP operating in Cluster-Mode.

```

# sanlun lun show -p all

          ONTAP Path: vs39:/vol/vol24_3_0/lun24_0
            LUN: 12
          LUN Size: 3g
      Host Device: /dev/rdisk/disk208
            Mode: C
            VG: /dev/vg01
      Multipath Policy: A/A
      Multipath Provider: Native
-----
host      vserver  /dev/dsk
path      path      filename
state     type      or hardware path
-----
up        primary  /dev/dsk/c3t1d4    fcd2  fc1_0    0
up        primary  /dev/dsk/c16t1d4   fcd2  fc1_1    0
up        primary  /dev/dsk/c130t1d4  fcd5  fc2_0    0
up        primary  /dev/dsk/c135t1d4  fcd5  fc2_1    0
up        secondary /dev/dsk/c8t1d4    fcd2  fc6_0    1
up        secondary /dev/dsk/c11t1d4   fcd2  fc6_1    1
up        secondary /dev/dsk/c13t1d4   fcd2  fc8_0    1
up        secondary /dev/dsk/c25t1d4   fcd2  fc8_1    1

```

The following output is of HP-UX 11iv3 using FCoE protocol.

```

# sanlun lun show -p all

          ONTAP Path: f3070-210-38:/vol/vol1/lun1
            LUN: 0
          LUN Size: 2g
      Controller CF State: Cluster Enabled
      Controller Partner: f3070-210-37
      Host Device: /dev/rdisk/disk5
            Mode: 7
            VG: /dev/vg01
      Multipath Policy: A/A
      Multipath Provider: Native
-----
host      controller /dev/dsk
path      path      filename
state     type      or hardware path
-----
up        primary  /dev/dsk/c1t0d0    fcoc4  1a      0
up        primary  /dev/dsk/c3t0d0    fcoc5  1b      0
up        secondary /dev/dsk/c5t0d0    fcoc4  1b      1
up        secondary /dev/dsk/c7t0d0    fcoc5  1a      1

```

(Veritas) Configuring LUNs for use with VxVM

If you are running Veritas Storage Foundation, you can use DMP for multipathing and VxVM to manage the LUNs. You can still use commands such as `sanlun` to display information about the LUN paths. You must perform several tasks when adding a LUN to a VxVM environment.

Steps

1. Confirm that you ran the `ioscan` and `ioinit` commands to ensure that the host detects the new LUN. Discover the LUNs on the host by using the steps outlined in the section [Discovering LUNs](#) on page 28.
2. Enable VxVM to detect the new LUN by entering the following command:

```
# vxctl enable
```

3. Scan the devices in the OS device tree by using following command:

```
# vxdisk scandisks
```

4. List the detected VxVM LUNs by entering the following command:

```
# vxdisk list
```

Example

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:LVM      -             -              LVM
disk_1          auto:LVM      -             -              LVM
fas30700_0     auto:none     -             -              online invalid
fas30700_1     auto:none     -             -              online invalid
```

The new disk devices are listed in the output of the `vxdisk list` command as online invalid disks. If you are using enclosure-based naming, the storage model is displayed in the `DEVICE` column. If you are using disk-based naming, the controller or disk name is displayed.

5. Initialize the disk for VxVM by entering the following command:

```
# /usr/lib/vxvm/bin/vxdisksetup -i device_name
```

Here *device_name* is the name listed for the LUN in the `DEVICE` column of the `vxdisk list` command output obtained in the previous step.

Example

```
# /usr/lib/vxvm/bin/vxdisksetup -i fas30700_0
# /usr/lib/vxvm/bin/vxdisksetup -i fas30700_1
```

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:LVM      -             -              LVM
disk_1          auto:LVM      -             -              LVM
fas30700_0     auto:cdsdisk -             -              online thinrclm
fas30700_1     auto:cdsdisk -             -              online thinrclm
```

6. Create a disk group for the LUN by entering the following command:

```
# vxdg init dg_name diskname=device_name
```

- *dg_name* is the name you assign to the disk group.
- *diskname* is the disk name you assign to the device you are adding to the disk group. The *diskname* represents the device.
- *device_name* is the controller name listed for the LUN in the `DEVICE` column of the `vxdisk list` command output.

Example

```
# /usr/sbin/vxdg init n_dg disk-1=fas30700_0
```

7. Add disks to the disk group by entering the following command:

```
# vxdg -g dg_name adddisk diskname=device_name
```

Example

```
# /usr/sbin/vxdg -g n_dg adddisk disk-2=fas30700_1
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:LVM      -             -              LVM
disk_1          auto:LVM      -             -              LVM
fas30700_0      auto:cdsdisk  disk-1        n_dg           online thinrclm
fas30700_1      auto:cdsdisk  disk-2        n_dg           online thinrclm
```

8. Create a logical volume by entering the following command:

```
# vxassist -g dg_name make vol_name size
```

- *dg_name* is the name of the disk group that you defined in Step 6.
- *vol_name* is the name you assign to the logical volume.
- *size* is the volume size.

Note: Size cannot be equal to or exceed the size of the LUN.

Example

```
# /usr/sbin/vxassist -g n_dg make nvol-1 1g
```

9. Create a file system on the volume by entering the following command:

```
# mkfs -F vxfs /dev/vx/rdisk/dg_name/vol_name
```

- *dg_name* is the name of the disk group that you defined previously.
- *vol_name* is the name you assigned to the logical volume you defined.

Example

```
# mkfs -F vxfs /dev/vx/rdisk/n_dg/nvol-1
version 9 layout
1048576 sectors, 1048576 blocks of size 1024, log size 16384 blocks
rcq size 1024 blocks
largefiles supported
```

10. Create a mount point by entering the following command:

```
# mkdir /mnt/dg_ntap01
```

11. Mount the new file system by entering the following command:

```
# mount logical_volume mount_point
```

Example

```
# mount /dev/vx/dsk/n_dg/nvol-1 /mnt/dg_ntap01
```

12. Verify the paths to the LUN by entering the following command:

```
# sanlun lun show -p all
```

Example

The following example shows output for the `sanlun` command.

The output you see varies depending on whether you are using HP-UX 11iv3 or HP-UX 11iv2.

HP-UX 11iv3 output

```
# sanlun lun show -p

                ONTAP Path: f3070-210-38:/vol/vol1/lun0
                  LUN: 0
                  LUN Size: 2g
    Controller CF State: Cluster Enabled
    Controller Partner: f3070-210-37
      Host Device: /dev/rdisk/disk10
        Mode: 7
      DMP NODE: fas30700_0
    Multipath Provider: Veritas
-----
```

host path state	controller path type	/dev/dsk filename or hardware path	host adapter	controller target port
up	primary	/dev/dsk/c8t0d0	fcd1	1b
up	primary	/dev/dsk/c6t0d0	fcd2	1a
up	secondary	/dev/dsk/c12t0d0	fcd1	1a
up	secondary	/dev/dsk/c10t0d0	fcd2	1b

HP-UX 11iv3 with FCoE LUN output

```
# sanlun lun show -p

                ONTAP Path: f3070-210-38:/vol/vol1/lun0
                  LUN: 0
                  LUN Size: 2g
    Controller CF State: Cluster Enabled
    Controller Partner: f3070-210-37
      Host Device: /dev/rdisk/disk10
        Mode: 7
      DMP NODE: fas30700_0
    Multipath Provider: Veritas
-----
```

host path state	controller path type	/dev/dsk filename or hardware path	host adapter	controller target port
up	primary	/dev/dsk/c8t0d0	fcoc4	1b
up	primary	/dev/dsk/c6t0d0	fcoc5	1a
up	secondary	/dev/dsk/c12t0d0	fcoc4	1a
up	secondary	/dev/dsk/c10t0d0	fcoc5	1b

HP-UX 11iv2 output

```
# sanlun lun show -p

                ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
                  LUN: 1
                  LUN Size: 100m
```

```

Controller CF State: Cluster Enabled
Controller Partner: f3070-210-38
Mode: 7
DMP NODE: FAS30700_0
Multipath Provider: Veritas
-----
host      controller /dev/dsk
path      path      filename
state     type      or hardware path      host      controller
-----
up        secondary /dev/dsk/c11t0d1      fcd1     1b
up        primary   /dev/dsk/c9t0d1       fcd1     1a
up        secondary /dev/dsk/c6t0d1       fcd0     1a
up        primary   /dev/dsk/c4t0d1       fcd0     1b

```

The sanlun utility

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

Displaying host LUN information with `sanlun`

You can use `sanlun` to display information about the LUNs connected to the host.

Steps

1. Ensure that you are logged in as root on the host.
2. Display LUN information by entering the following command:

```
# sanlun lun show
```

This command has the following basic formats:

```
sanlun lun show [-v] [-d <host_device_filename> |
    all |
    <controller/vserver_name> |
    <controller/vserver_name>:<path_name>]
sanlun lun show -wwpn [ <target_wwpn> |
    <colon(:)_separated_target_wwpn>]
sanlun lun show -p [-v] [ all |
    <controller/vserver_name> |
    <controller/vserver_name>:<path_name>]
sanlun lun show -b [ all |
    <controller/vserver_name> |
    <controller/vserver_name>:<path_name>]
sanlun fcp show adapter [ -c | [ -v ] [<adapter_name> | all ]]
sanlun version
sanlun [ lun | fcp ] help
```

Note: You must use the `-p` version of the command to see the information about the primary and secondary paths. You cannot use the `-d` option if you use the `-p` option.

- `-p` displays multipathing information.
- `-wwpn` displays all devices on the Data ONTAP controller or Vserver with the target FC port `<wwpn>`.
- `all` displays information for all LUNs that are currently attached to the host.
- `-b` prints a brief listing of the LUNs.
- `-d` specifies the device special file on the host (*host_device_filename*). For example for HP-UX 11iv3, the output might be `/dev/rdisk/disk56`.

controller_name or *vserver_name* is the name of the target storage system.

path_name is the path to the LUN on the storage system.

- `-v` produces verbose output.

If you enter `sanlun lun show`, `sanlun lun show -p`, or `sanlun lun show -v` without any parameters, the utility responds as if you had included all parameters.

The options you enter depend on the information you want to view. The following command lines illustrate the different ways you might want to use the `sanlun` command.

- Display a list of the paths and multipathing information associated with the LUN by entering the following command:

```
# sanlun lun show -p
```

- Display the summary listing of the storage system LUN associated with the host device `/dev/rdisk/disk223` by entering the following command:

```
# sanlun lun show -d /dev/rdisk/disk223
```

- Display verbose output for all LUNs currently available on the host by entering the following command:

```
# sanlun lun show -v all
```

- Display a summary listing of all LUNs available to the host from the storage system named `controllerA` by entering the following command:

```
# sanlun lun show controllerA
```

- Display a summary listing of all LUNs available to the host from the storage system named `vserverA` by entering the following command:

```
# sanlun lun show vserverA
```

- Display a list of paths between the host and a specific LUN on the storage system `controllerA` by entering the following command:

```
# sanlun lun show controllerA:path_name
```

LUNs that were created on the storage system but not discovered by the host are not displayed. Also, the device special files must be created before the LUN can be displayed.

- Display a list of paths between the host and a specific LUN on the storage system for `vserverA` by entering the following command:

```
# sanlun lun show vserverA:path_name
```

Example

The examples that follow show the sample output when you use the `sanlun lun show` command with the `-p` option.

Example using LVM

On a system using LVM, you see output similar to the following if you enter the `sanlun lun show` command with the `-p` option.

```
# sanlun lun show -p
```

```

ONTAP Path: vs39:/vol/vol24_3_0/lun24_0
LUN: 12
LUN Size: 3g
Host Device: /dev/rdisk/disk208
Mode: C
VG: /dev/vg01
Multipath Policy: A/A
Multipath Provider: Native
-----
host      vsserver  /dev/dsk
path      path      filename
state    type      or hardware path      host      vsserver  HP A/A
                                adapter  LIF      path failover
                                adapter  LIF      priority
-----
up        primary  /dev/dsk/c3t1d4      fcd2     fc1_0     0
up        primary  /dev/dsk/c16t1d4     fcd2     fc1_1     0
up        primary  /dev/dsk/c130t1d4    fcd5     fc2_0     0
up        primary  /dev/dsk/c135t1d4    fcd5     fc2_1     0
up        secondary /dev/dsk/c8t1d4      fcd2     fc6_0     1
up        secondary /dev/dsk/c11t1d4     fcd2     fc6_1     1
up        secondary /dev/dsk/c13t1d4     fcd2     fc8_0     1
up        secondary /dev/dsk/c25t1d4     fcd2     fc8_1     1

```

Example using DMP:

On a system using DMP, you see output similar to the following if you enter the `sanlun lun show` command.

```

# sanlun lun show -p
ONTAP Path: f3070-210-38:/vol/vol1/lun0
LUN: 0
LUN Size: 2g
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-37
Host Device: /dev/rdisk/disk10
Mode: 7
DMP NODE: N52000_94
Multipath Provider: Veritas
-----
host      controller /dev/dsk
path      path      filename
state    type      or hardware path      host      controller
                                adapter  target
                                adapter  port
-----
up        primary  /dev/dsk/c8t0d0      fcd1     1b
up        primary  /dev/dsk/c6t0d0      fcd2     1a
up        secondary /dev/dsk/c12t0d0     fcd1     1a
up        secondary /dev/dsk/c10t0d0     fcd2     1b

```

Example using LVM with ALUA while legacy mode is disabled

On an HP-UX 11iv3 system using LVM with ALUA enabled and legacy I/O nodes and legacy DSFs disabled on the host, `sanlun` returns the following output:

```

# sanlun lun show -p
ONTAP Path: vs39:/vol/vol24_3_0/lun24_0
LUN: 12
LUN Size: 3g
Host Device: /dev/rdisk/disk208
Mode: C
VG: /dev/vg01
Multipath Policy: A/A
Multipath Provider: Native
-----
host      vsserver  /dev/dsk
path      path      filename
state    type      or hardware path      host      vsserver  HP A/A
                                adapter  LIF      path failover
                                adapter  LIF      priority
-----
up        primary  0/3/0/0/0/0.0x210000a098113766.0x400c000000000000
                                fcd2     fc1_0     0
up        primary  0/3/0/0/0/0.0x220000a098113766.0x400c000000000000
                                fcd2     fc1_1     0
up        primary  0/7/0/0/0/0.0x230000a098113766.0x400c000000000000

```


The following examples show the type of output you see when you use different options of the `sanlun fcp show adapter` command.

```
# sanlun fcp show adapter

fcd0          WWPN:50060b000060ea16
fcd1          WWPN:50060b000060ea18
```

Adding the `-c` option, so that you enter the following command:

```
# sanlun fcp show adapter -c
```

```
# sanlun fcp show adapter -c

Enter the following command on controllers (running in 7-mode only)
to create an initiator group for this system:
igroup create -f -t hpux "hpux_24" 50060b000060ea16 50060b000060ea18
50014380029cad36
```

Adding the `-v` option, so that you enter the following command:

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

```
# sanlun fcp show adapter -v

adapter name:      fcd2
WWPN:              50014380029cad36
WWNN:              50014380029cad37
driver name:       fcd
model:             AK344A or AH400A
model description: HP 8Gb Single Channel PCI-e 2.0 FC HBA
serial number:     MY5915201W
hardware version:  2
driver version:    @(#) fcd B.11.31.1109 May 23 2011
firmware version:  5.4.4
Number of ports:  1
port type:         Fabric
port state:        Operational
supported speed:   8 GBit/sec
negotiated speed:  8 GBit/sec
OS device name:    /dev/fcd2

adapter name:      fcd5
WWPN:              500143800169831a
WWNN:              500143800169831b
driver name:       fcd
model:             AK344A or AH400A
model description: HP 8Gb Single Channel PCI-e 2.0 FC HBA
serial number:     MY5837301N
hardware version:  2
driver version:    @(#) fcd B.11.31.1109 May 23 2011
firmware version:  5.4.4
Number of ports:  1
port type:         Fabric
port state:        Operational
supported speed:   8 GBit/sec
```

```
negotiated speed: 8 GBit/sec
OS device name: /dev/fcd5
```

Output of `sanlun fcp show adapter -v` command when FCoE adapter is present on host.

```
# sanlun fcp show adapter -v
adapter name: fcoc4
WWPN: 1000d8d385d52aa1
WWNN: 2000d8d385d52aa1
driver name: fcoc
model: 580153-001
model description: HP NC551m Converged Network Adapter
serial number: THC01704WC
hardware version: 2
driver version: @(#) FCOC: PCIe Fibre Channel driver (FibrChan1-03),
B.11.31.1103, Dec 4 2010, FCOC_IFC (4,1)
firmware version: 1.10R0 SLI-2 (2.702.485.4)
Number of ports: 1 of 2
port type: Fabric
port state: Operational
supported speed: 10 GBit/sec
negotiated speed: 10 GBit/sec
OS device name: /dev/fcoc4

adapter name: fcoc5
WWPN: 1000d8d385d52aa5
WWNN: 2000d8d385d52aa5
driver name: fcoc
model: 580153-001
model description: HP NC551m Converged Network Adapter
serial number: THC01704WC
hardware version: 2
driver version: @(#) FCOC: PCIe Fibre Channel driver (FibrChan1-03),
B.11.31.1103, Dec 4 2010, FCOC_IFC (4,1)
firmware version: 1.10R0 SLI-2 (2.702.485.4)
Number of ports: 2 of 2
port type: Fabric
port state: Operational
supported speed: 10 GBit/sec
negotiated speed: 10 GBit/sec
OS device name: /dev/fcoc5
```

How sanlun displays the multipath provider

You can use the `sanlun` command to display multipathing information about paths to LUNs connected to the HP-UX host.

You can use either the VxVM or the LVM to manage volumes. The information that `sanlun` command displays differs based the volume manager used or whether you use a volume manager.

- When the LUN is controlled by VxVM, the multipath provider is displayed as Veritas.
- If the LUN is controlled by the HP-UX LVM, the output displays the multipath provider as Native and the policy as either A/A or A/P.

- If the LUN is not controlled by a volume manager, the multipath provider is none.

(Veritas) Veritas Dynamic Multipathing configuration

If you are using the Veritas Volume Manager with Dynamic Multipathing (DMP) for multipathing support on an HP-UX system, you must ensure that the VxVM software, Veritas patch bundle, and Array Support Library (ASL) and Array Policy Module (APM) are installed.

(Veritas) What the ASL is

The ASL is a Data ONTAP-qualified library that provides information about storage array attributes and multipathing configurations to the Device Discovery Layer (DDL) and Veritas Dynamic Multipathing (DMP) components of Veritas Volume Manager (VxVM).

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

Note: You cannot use storage systems simultaneously as JBOD 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.

(Veritas) ASL array type

The ASL reports information about the multipathing configuration to the DDL as an Active/Active (A/A), ALUA, or an Active/Passive Concurrent (A/P-C) disk array type.

- **Active/Active (A/A)**
There are multiple active paths to a storage system, and simultaneous I/O is supported on each path. If a path fails, I/O is distributed across the remaining paths.
- **Active/Passive Concurrent (A/P-C)**
An A/P-C array is a variant of the A/P 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.
- **ALUA**
A LUN in an ALUA-enabled array can be accessed through both controllers by using optimized and non-optimized paths. The array notifies the host of path options, their current state, and state changes. Using this information, the host can determine which paths are optimized. Failover to the non-optimized path occurs only if all the optimized paths fail.

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

(Veritas) Information about ASL error messages

Normally, the ASL works silently and seamlessly with the VxVM DDL. If an error, malfunction, or misconfiguration occurs, messages from the library are logged to the console using the host's logging facility. The ASL error messages have different levels of severity and importance.

If you receive one of these messages, call Symantec Technical Support for help. The following table lists the importance and severity of these messages.

Message severity	Definition
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.
Warning	Indicates that an UNCLAIMED status is being returned. Unless claimed by a subsequent ASL, dynamic multipathing is disabled. No error is being returned but the device (LUN) might not function as expected.
Info	Indicates that a CLAIMED status is being returned. The device functions fully with Veritas DMP enabled, but the results seen by the user might be other than what is expected. For example, the enclosure name might change.

(Veritas) What the APM is

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

The Symantec APM for IBM N series storage arrays is customized to optimize I/O error handling and failover path selection for the N series environment.

After the ASL discovers the storage array as an N series array, the ASL instructs DMP to use the N series specific APM to handle I/O error processing and path failures for the N series storage array.

(Veritas) How to get the ASL and APM

The ASL and APM are available from the Symantec website. They are not included with the Host Utilities.

To determine which versions of the ASL and APM you need for your version of the host operating system, check the [IBM N series interoperability matrix website](#). This information is updated frequently. When you know which version you need, go to the Symantec website and download the ASL and APM.

Note: Because the ASL and APM are Symantec (Veritas) products, Symantec provides technical support if you encounter a problem using them.

Note: From Veritas Storage Foundation 5.1 onwards, the ASL and APM are included in the Veritas Storage Foundation product.

For Veritas Storage Foundation 5.0, the Symantec TechNote download file contains the software packages for both the ASL and the APM. You must extract the software packages and then install each one separately as described in the TechNote.

Information about getting the Symantec TechNote for the ASL and APM is provided on the IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

(Veritas) Installing the ASL and APM software

If you are using Veritas Storage Foundation for multipathing, you should install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM) for storage systems.

Before you begin

- Verify that your configuration meets the system requirements.
- Obtain the ASL and APM software.

The ASL and APM are not distributed with the Host Utilities software. You can obtain the ASL and APM from the Symantec website. See the appropriate interoperability matrix at IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

You should download this software before you start the installation.

About this task

Only one version of the ASL and APM package can be installed on the host at any given time.

Steps

1. Log in to the HP-UX host.

2. If you already have the storage configured as JBOD in your VxVM configuration, remove the JBOD support for storage by entering the following command:

```
vxddladm rmjbod vid=ONTAP
```

3. Install the ASL and APM according to the instructions provided by Symantec.

Example

- a. Create a directory to store the VRTSaslapm package.

```
# mkdir /tmp/aslapm
```

```
# cd /tmp/aslapm
```

- b. Copy VRTSaslapm_HPUX_6.0.000.100.tar.gz in /tmp/aslapm directory and verify the cksum.

```
# cksum VRTSaslapm_HPUX_6.0.000.100.tar.gz
```

```
3250523557 3967977 VRTSaslapm_HPUX_6.0.000.100.tar.gz
```

- c. Uncompress the file and extract the packages.

```
# gunzip VRTSaslapm_HPUX_6.0.000.100.tar.gz
```

```
# tar -xvf VRTSaslapm_HPUX_6.0.000.100.tar
```

- d. If VRTSaslapm package is currently NOT installed on your system, then perform an initial install.

```
# swinstall -s `pwd` VRTSaslapm
```

- e. After you install VRTSaslapm package, you must execute:

```
# vxdctl enable
```

4. If your host is connected to a storage system, verify the installation by following these steps:

- a. Run the following command:

```
# vxdmpadm listenclosure all
```

The output shows the model name of the storage device if you are using enclosure- based naming with VxVM.

Example

The `vxdmpadm listenclosure all` command shows the Enclosure Type as N5600 in this example.

```
# vxdmpadm listenclosure all
ENCLR_NAME      ENCLR_TYPE      ENCLR_SNO      STATUS      ARRAY_TYPE      LUN_COUNT
=====
disk            Disk            DISKS          CONNECTED   Disk            2
N5600          N5600          3079846       CONNECTED   ALUA            2
```

Related information

IBM N series interoperability matrix website: www.ibm.com/systems/storage/network/interphome.html

(Veritas) Uninstalling the ASL and APM

You uninstall the ASL and APM using the `swremove` command. Before you uninstall these packages, you should make sure your system is in the correct state.

Perform the following tasks before you uninstall the ASL and APM:

- Stop all I/O to LUNs.
- Deport the disk group.

Note: In a Storage Foundation RAC, you must also stop clustering on a node before you remove the ASL and APM.

Example of uninstalling the ASL and APM

The following is an example of uninstalling the ASL and the APM when you have Veritas Storage Foundation 6.0.

Note: If you were actually doing this uninstall, your output would vary slightly based on your system setup. Do not expect to get identical output on your system.

```
bash-2.05# swremove VRTSaslapm

===== 04/04/12 16:19:51 IST BEGIN swremove SESSION
         (non-interactive) (jobid=hpux_20-0029)

         * Session started for user "root@hpux_20".

         * Beginning Selection
         * Target connection succeeded for "hpux_20:/".
NOTE:    The software specified contains a kernel fileset or a
         dynamic_module fileset. The kernel will be modified, and if
         necessary the system will be rebooted.
         * Software selections:
           VRTSaslapm.VXASLAPM-KRN,l=/,r=6.0.000.000,a=HP-UX_B.11.31_IA/PA,v=Symantec,fr=6.0.000.000,fa=HP-
UX_B.11.31_PA
           VRTSaslapm.VXASLAPM-RUN,l=/,r=6.0.000.000,a=HP-UX_B.11.31_IA/PA,v=Symantec,fr=6.0.000.000,fa=HP-
UX_B.11.31_PA
         * Selection succeeded.

         * Beginning Analysis
         * Session selections have been saved in the file
           "/.sw/sessions/swremove.last".
         * The analysis phase succeeded for "hpux_20:/".
         * Analysis succeeded.

         * Beginning Execution
         * The execution phase succeeded for "hpux_20:/".
         * Execution succeeded.

NOTE:    More information may be found in the agent logfile using the
         command "swjob -a log hpux_20-0029 @ hpux_20:/".

===== 04/04/12 16:21:20 IST END swremove SESSION (non-interactive)
         (jobid=hpux_20-0029)
```

(Veritas) Upgrading the ASL and APM

If you are using DMP with Veritas Storage Foundation 5.0 or later, you must install the ASL and the APM. If the ASL and APM are already installed and you need to upgrade them to a newer version,

you must first remove the existing versions. Then you can obtain and install the new ASL and APM software packages.

Steps

1. If you currently have the ASL installed, you should check its version to determine whether you need to update it. Use the Veritas `vxddladm listversion` command to determine the ASL version.

The `vxddladm listversion` command generates the following output.

```
# vxddladm listversion
LIB_NAME                               ASL_VERSION           Min. VXVM version
=====
libvxlsialua.sl                        vm-6.0-rev-2         6.0
libvxmsa2k.sl                          vm-6.0-rev-1         6.0
libvxmsa2kfc_sa.sl                    vm-6.0-rev-1         6.0
libvxnetapp.sl                        vm-6.0-rev-1       6.0
```

This output has been truncated to make the document easier to read.

2. Determine the ASL and APM packages name by executing the `swlist` command.

The `swlist` command generates the following output.

Note: The output might change depending on the version of ASL you are using. This example is of Veritas Storage Foundation 6.0.

```
# swlist|grep -i asl
VRTSaslapm                               6.0.000.000          Array Support Libraries and Array Policy Modules for
Veritas Volume Manager

# swlist|grep -i apm
VRTSaslapm                               6.0.000.000          Array Support Libraries and Array Policy Modules for
Veritas Volume Manager
```

3. Uninstall the already present ASL and APM versions by using the `swremove` command.
4. Install the ASL and APM version to which you want to upgrade using the `swinstall` command.

(Veritas, HP-UX 11iv3) Disabling Native MPIO ALUA

When you are using Veritas DMP and HP-UX 11iv3, you must use the ASL to ensure that DMP functions properly in the Veritas Storage Foundation stack. This also means that you must disable Native MPIO ALUA on Veritas LUNs.

Before you begin

Halt I/O before you start. You should not run I/O while you are making this change.

About this task

If you do not disable Native MPIO ALUA on Veritas LUNs, the `sanlun lun show -p` command output does not display the DMP node names for the devices.

There are different ways to disable Native MPIO ALUA. The following steps disable it at the LUN level by using the `scsimgr` command to set the `alua_enabled` attribute for each LUN to 0 (false). You must also make the attribute persistent across host reboots.

Steps

1. Determine whether the `alua_enabled` attribute has already been set to false and made persistent by entering the command:

```
# scsimgr -p get_attr all_lun -a device_file -a alua_enabled
```

Example

In this example, the attribute has not been set to false or made persistent.

Note: `disk460` is used as the sample LUN in the examples.

```
# scsimgr get_attr -D /dev/rdisk/disk460 -a alua_enabled

SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

name = alua_enabled
current = true
default = true
saved =
```

2. Set the attribute to false by entering the command:

```
# scsimgr set_attr -D /dev/rdisk/disk460 -a alua_enabled=0
```

3. Make sure the attribute persists across reboots by entering the command:

```
# scsimgr save_attr -D /dev/rdisk/disk460 -a alua_enabled=0
```

Note: These commands take effect immediately. You do not need to reboot the host.

4. Verify that attribute has been changed to false and made persistent across host reboots by entering the command:

```
# scsimgr get_attr -D /dev/rdisk/disk460 -a alua_enabled
```

Example

This example shows the updated settings for the attribute.

```
# scsimgr get_attr -D /dev/rdisk/disk460 -a alua_enabled

SCSI ATTRIBUTES FOR LUN : /dev/rdisk/disk460

name = alua_enabled
current = false
```

```
default = true
saved = false
```

Related concepts

[HP-UX configurations that support ALUA](#) on page 16

Related information

[How to Disable HP-UX 11iv3 Native Multi-Pathing ALUA mode for Storage Foundation 5.0 and 5.0.1](#)

(Veritas) The Veritas DMP restore daemon requirements

You must set the Veritas restore daemon values for the restore policy and the polling interval to the Host Utilities recommended values. These settings determine how frequently the Veritas daemon checks paths between the host and the storage system.

At the time this document was produced, the Host Utilities recommended values were as follows:

- restore daemon interval: 60 seconds
- restore daemon policy: check_disabled
- vxpfto: 30 seconds (default)

Check the *Release Notes* to see if these values have changed since this document was produced.

(Veritas) The Veritas DMP I/O policy

The command you use to view the I/O policy on the enclosure varies depending on which version of HP-UX you are running.

To get information about the I/O policy when you are using either HP-UX 11iv3 or HP-UX 11iv2, use the command `vxddmpadm getattr enclosure storage_system iopolicy`. This command displays output similar to the following:

```
# /usr/sbin/vxddmpadm getattr enclosure fas30700 iopolicy
ENCLR_NAME      DEFAULT          CURRENT
=====
fas30700        MinimumQ        MinimumQ
```

When you are using HP-UX 11iv3 for Agile DSF, use the command `vxddmpadm list dmpnode` to display information about the I/O policy. This command displays output similar to the following:

```
# /usr/sbin/vxddmpadm list dmpnode dmpnodename=disk10
dmpdev          = fas30700_0
state           = enabled
enclosure       = fas30700
cab-sno         = 3079846
asl             = libvxontap.sl
vid             = ONTAP
```

```

pid = LUN
array-name = FAS3070
array-type = ALUA
iopolicy = MinimumQ
avid = -
lun-sno = 1kZtf]A1M/gH
udid = ONTAP%5FLUN%5F3079846%5F1kZtf%5DA1M%2FgH
dev-attr = tprclm
lun_type = std
scsi3_vpd = 60A98000316B5A74665D41314D2F6748
num_paths = 4
###path = name state type transport ctrl hwpath aprotID aprotWWN attr
path = c8t0d0 enabled(a) primary FC c8 0/4/2/0.0x500a09848739357d 2-4 - -
path = c12t0d0 no_license secondary FC c12 0/4/2/0.0x500a09839739357d 3-3 - -
path = c10t0d0 no_license secondary FC c10 0/4/2/1.0x500a09849739357d 3-4 - -
path = c6t0d0 no_license primary FC c6 0/4/2/1.0x500a09838739357d 2-3 - -

```

(Veritas) Displaying multipathing information using sanlun

You can use the Host Utilities' `sanlun` utility to display information about the array type and paths to LUNs on the storage system in Veritas DMP environments using ASL and APM.

About this task

When ASL is installed the LUN displays multipath provider as Veritas.

Step

1. On the host, enter the following command:

```
# sanlun lun show -p all
```

The `sanlun` utility displays path information for each LUN; however, it only displays the native multipathing policy. To see the multipathing policy for other vendors, you must use vendor-specific commands.

(Veritas) Displaying available paths using VxVM

You can use VxVM to display information about available paths to a LUN.

Steps

1. View all the devices by entering:

```
# vxdisk list
```

The VxVM management interface displays the `vxdisk` device, type, disk, group, and status. It also shows which disks are managed by VxVM.

The following example shows the type of output you see when you enter the `vxdisk list` command.

```

# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
disk_0          auto:LVM      -             -             LVM

```

```
disk_1      auto:LVM      -      -      LVM
fas30700_0  auto:cdsdisk  disk-1  n_dg   online thinrclm
fas30700_1  auto:cdsdisk  -      -      online thinrclm
```

This output has been truncated to make the document easier to read.

2. On the host console, display the path information for the device you want by entering:

```
# vxddmpadm getsubpaths dmpnodename=device
```

where *device* is the name listed under the output of the `vxddisk list` command.

The following example shows the type of output you see when you enter this command.

```
# vxddmpadm getsubpaths dmpnodename=fas30700_0
NAME          STATE [A]    PATH-TYPE [M]  CTLR-NAME    ENCLR-TYPE   ENCLR-NAME    ATTRS
-----
c10t0d0      ENABLED      SECONDARY      c10          FAS3070     fas30700      -
c12t0d0      ENABLED      SECONDARY      c12          FAS3070     fas30700      -
c6t0d0       ENABLED (A)  PRIMARY        c6           FAS3070     fas30700      -
c8t0d0       ENABLED (A)  PRIMARY        c8           FAS3070     fas30700      -
```

3. To obtain path information for a host HBA, enter:

```
# vxddmpadm getsubpaths ctlr=controller_name
```

controller_name is the controller displayed under CTLR-NAME in the output of the `vxddmpadm getsubpaths dmpnodename` command you entered in Step 2.

The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to.

The following example displays the type of output you should see.

```
# vxddmpadm getsubpaths ctlr=c8
NAME          STATE [A]    PATH-TYPE [M]  DMPNODENAME  ENCLR-TYPE   ENCLR-NAME    ATTRS
-----
c8t0d0      ENABLED (A)  PRIMARY        fas30700_0   FAS3070     fas30700      -
c8t0d1      ENABLED      SECONDARY      fas30700_1   FAS3070     fas30700      -
```

This output has been truncated to make the document easier to read.

(Veritas, HP-UX 11iV3) Thin provisioning and space reclamation

When you are running HP-UX 11iV3 and Veritas Storage Foundation 5.0.1 RP2 or later, you can use the VxFS space reclamation feature. For more information about configurations that support this feature, see IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

The space reclamation feature requires that you have the appropriate ASL installed.

Contact Symantec Support to get this ASL.

You can use the `vxddisk reclaim` command to reclaim space at the following levels:

- DMP node
- DMP enclosure
- Veritas disk group
- File system

The following examples show how you can reclaim space at these levels.

This command reclaims space at the DMP node level.

```
vxdisk reclaim fas30700_12
Reclaiming thin storage on:
Disk fas30700_12 : Done
```

This command reclaims space at the DMP enclosure level.

```
vxdisk reclaim fas30700
Reclaiming thin storage on:
Disk fas30700_12 : Done!
Disk fas30700_13 : Done!
Disk fas30700_14 : Done!
Disk fas30700_15 : Done
```

This command reclaims space at the Veritas disk group level.

```
vxdisk reclaim dg1
Reclaiming thin storage on:
Disk fas30700_15 : Done!
Disk fas30700_12 : Done!
Disk fas30700_13 : Done!
Disk fas30700_14 : Done
```

This command reclaims space at the file system level.

```
# fsadm -F vxfs -R /mnt/qa/n_vg/nvol-1
```

(HP-UX) Setting up SAN boot LUN on HP-UX

You can set up a SAN boot LUN to work in an HP-UX environment that is using the FC/FCoE protocol.

Before you begin

Verify that your system configuration supports SAN boot LUNs. See the IBM N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 6).

Note: At the time this document was produced, HP-UX environment only supported SAN boot with the FC/FCoE protocol.

Steps

1. Ensure that you have zoned the HP-UX host and the controllers correctly.
2. Create a LUN on the storage system and map it to the host.

This LUN will be the SAN boot LUN.

You should ensure the following:

- The LUN 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.
 - The LUN is visible to the host during the boot process.
3. Boot from HP-UX operating system DVD media or ignite server.
 4. From the Installation Process console, select the Install HP-UX option.

```
Welcome to the HP-UX installation/recovery process!
```

```
Use the <tab> key to navigate between fields, and the arrow keys
within fields. Use the <return/enter> key to select an item.
Use the <return/enter> or <space-bar> to pop-up a choices list. If the
menus are not clear, select the "Help" item for more information.
```

```
Hardware Summary:          System Model: ia64 hp server rx3600
+-----+-----+-----+-----+
| Disks: 3 ( 286.7GB) | Floppies: 0 | LAN cards: 2 | [ Scan Again ] |
| CD/DVDs: 1 | Tapes: 0 | Memory: 16353Mb | |
| Graphics Ports: 1 | IO Buses: 7 | CPUs: 4 | [ H/W Details ] |
+-----+-----+-----+-----+

          [      Install HP-UX      ]

          [ Run an Expert Recovery Shell ]

          [      Advanced Options      ]

[ Reboot ] [ Help ]
```


The installation process has begun.

```

+-----+
+                                     /opt/ignite/bin/itool ()
+-----+
+-----+
+ Basic  || Software || System || File System || Advanced |
+-----+
+-----+
+ Configurations: [ HP-UX B.11.31 Default    ->] [ Description... ]
+ Environments:   [ HP-UX Data Center Operation ->] (HP-UX B.11.31)
+ [ Root Disk ]  ONTAP_LUN, 0/3/0/0/0/0.0x500a0983973935>
+ File System:    [ Logical Volume Manager (LVM) with VxFS ->]
+ [ Root Swap (MB)... ] 8192      Physical Memory (RAM) = 16353 MB
+ [ Languages... ] English          [ Keyboards... ] [ Additional... ]
+-----+
+ [ Show Summary... ]                                     [ Reset Configuration ]
+-----+
+ [ Go! ] [ Cancel ] [ Help ]
+-----+

```

10. After the OS installation is complete, install the Host Utilities.

Note: For more information, refer to the HP-UX installation document from HP.

Troubleshooting

If you encounter a problem while running the Host Utilities, here are some tips and troubleshooting suggestions that might help you resolve the issue.

This chapter contains the following information:

- Best practices, such as checking the *Release Notes* to see if any information has changed.
- Suggestions for checking your system.
- Information about possible problems and how to handle them.
- Diagnostic tools that you can use to gather information about your system.

Host setup checklist

The following checklist pertains to the host setup. You should confirm that the following statements are true.

- Output from the `sanlun fcp show adapter -v` command shows that the HBAs are operational.
- All the LUN have been discovered by the host.
- The software states reported by the `ioscan` command for LUNs are CLAIMED (that is, in a good state).
- The `sanlun lun show -p` command does not report problems such as the following:
 - Downed paths
 - Missing paths
 - Out-of-order paths
 - The cluster in a takeover state
- No uncertain messages are 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.
- No other volume group configuration problems are reported. You can use the `vgdisplay -v` command to display information about volume group configurations.

Storage system setup checklist

You should confirm that the following statements are true for your storage system:

- FC/iSCSI is running and all the target ports are in a good state.
 - Note:** If an HP-UX FC driver is not in use, it logs off the target port. For this reason, the `fcshow initiators` command might not display all the host drivers as connected even though they are.
- The igroups must have the `ostype` attribute as `hpux`.
- For LUNs, the following statements are true:
 - The `ostype` attribute is `hpux`.
 - LUNs are mapped to host igroups.
 - Output from the `sanlun lun show -p` command lists the LUN state as good.

Related information

IBM N series interoperability matrix website: www.ibm.com/systems/storage/network/interphome.html

Connectivity problems when using LVM

On HP-UX systems using LVM, applications do not generally produce I/O errors. Instead, I/O retries continue uninterrupted if connectivity is lost.

If I/O retries appear to be continuing without succeeding, check the following areas for connectivity problems:

- System setup
- Connectivity between the storage system and the FC switch
- FC switch setup
- FC connectivity to the host
- Host setup

Migrating a configuration from non-ALUA to ALUA without a host reboot

You can migrate from non-ALUA to ALUA without rebooting the host when using HP-UX 11iv3 native multipathing.

Steps

1. On the controller, capture the existing LUN statistics information by entering the following command:

```
lun stats -o
```

2. Reset the statistics to zero by entering the following command:

```
lun stats -z
```

3. Run `igroup set` command:

```
igroup set igroup_name ALUA
```

4. Capture the new LUN statistics information by entering the following command:

```
lun stats -o
```

5. On the host, perform `ioscan`

6. Disable port `fcd0` by entering the following command and wait for 30 seconds:

```
fcmsutil /dev/fcd0 disable
```

7. Enable port `fcd0` by entering the following command and wait for 30 seconds:

```
fcmsutil /dev/fcd0 enable
```

8. Disable port on `fcd1` by entering the following command and wait for 30 seconds:

```
fcmsutil /dev/fcd1 disable
```

9. Enable port `fcd1` by entering the following command and wait for 30 seconds:

```
fcmsutil /dev/fcd1 enable
```

10. Perform `ioscan`

11. Repeat Step 6 through Step 12 for all hosts with WWPNs of `igroup`.

12. Verify that I/O is going through the active/optimized path as expected.

(Veritas) Application I/O failures on Veritas VxVM

Application I/O failures might sometimes occur during storage system operations, such as takeover or giveback. You can perform several checks to troubleshoot the failure.

- Verify that the vxpfto setting is set to the values specified in this guide.
- Verify that the DMP restore daemon settings are set to the values specified in this guide.

(Veritas) Enclosure-based naming is not working correctly

Verify the following when the enclosure-based naming on your system is not functioning as expected.

- The ASL is loaded correctly.
- You changed the disk-naming scheme to enclosure-based naming using the `vxdiskadm` utility.

Problems with volume group multipathing

When you are using HP-UX 11iv2 with PV-Links, you can use the `ontap_config_paths` utility to correct problems with volume group multipathing.

You can correct volume group multipathing by perform the following tasks using the `ontap_config_paths` utility:

- Set up alternative 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.

The following example shows how `ontap_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 example is truncated. Part of the `ontap_config_paths` output was removed.

```
# sanlun lun show -p
```

```

                ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
                LUN: 1
                LUN Size: 100m
    Controller CF State: Cluster Enabled
    Controller Partner: f3070-210-38
                Mode: 7
                VG: /dev/test_vg
    Multipath Policy: A/P
    Multipath Provider: Native

```

```
-----
host          controller /dev/dsk          controller  PVlinks
```

```

path      path      filename      host      target      path failover
state     type       or hardware  adapter  port        priority
-----
up        primary   /dev/dsk/c4t0d1  fcd0     1b          0
up        secondary /dev/dsk/c11t0d1 fcd1     1b          not in VG
up        primary   /dev/dsk/c9t0d1  fcd1     1a          not in VG
up        secondary /dev/dsk/c6t0d1  fcd0     1a          not in VG

```

ontap_config_paths

Getting information from sanlun...

Moving VG path to the end of the alternate paths with:

```

vgreduce /dev/test_vg /dev/dsk/c6t0d1
vgextend /dev/test_vg /dev/dsk/c6t0d1

```

Device file path "/dev/dsk/c6t0d1" is an primary link.
Removing primary link and switching to an alternate link.
Volume group "/dev/test_vg" has been successfully reduced.
Volume Group configuration for /dev/test_vg has been saved in /etc/lvmconf/
test_vg.conf

Current path "/dev/dsk/c4t0d1" is an alternate link, skip.
Current path "/dev/dsk/c11t0d1" is an alternate link, skip.
Volume group "/dev/test_vg" has been successfully extended.
Volume Group configuration for /dev/test_vg has been saved in /etc/lvmconf/
test_vg.conf

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.

Several things occur during a takeover:

- All the primary paths from the host that go through the failed storage system to its LUNs become unresponsive. If you run the `ioscan` command, the output displays a S/W State of `NO_HW`, which means unresponsive paths.
- Applications do not see I/O errors from the system.
- It might take PV-Links 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 this value, enter `pvdisplay -v`. To change this value, enter `pvchange`.
- The host's `syslog` file might contain timeout messages.

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

Events that occur if a host reboots during a takeover

When a reboot occurs, you must rediscover the primary paths to the volume groups.

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 might lose its secondary paths also.
- The volume groups deactivate the primary paths.
- The volume groups remain active but are aware of the secondary paths only. These paths go 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 secondary paths, not the primary paths.

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.

Recovering from a host reboot during a takeover

In HP-UX 11iv2, if the host reboot occurs while the cluster is in a takeover state, it does not automatically rediscover the paths to LUNs after giveback.

Steps

1. Ensure that controller comes to optimal state.
2. On the host, detect the primary paths by entering the `ioscan` command.

Example

The following sample output from the `ioscan` command shows all paths to the LUN:

```
# 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 ONTAP LUN
/dev/dsk/c24t0d0 /dev/rdisk/c24t0d0
disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE ONTAP LUN
/dev/dsk/c26t0d0 /dev/rdisk/c26t0d0
disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE ONTAP LUN
/dev/dsk/c46t0d0 /dev/rdisk/c46t0d0
disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE ONTAP LUN
/dev/dsk/c44t0d0 /dev/rdisk/c44t0d0
```

3. On the host, ensure that the host detects all paths to LUNs by entering the following command:

```
# sanlun lun show -p all
```

Example

The following example shows partial `sanlun` command output after you use the `ioscan` command to detect the recovered paths.

Note: The PV-Links priority is not correct. A secondary path is listed as the first path in the PV-Links order. The `sanlun` command shows all recovered paths, but the paths are not shown in the correct multipathing order.

```
# sanlun lun show -p
ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
LUN: 1
LUN Size: 100m
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-38
Mode: 7
VG: /dev/test_vg
Multipath Policy: A/P
Multipath Provider: Native
-----
host controller /dev/dsk controller PVlinks
path path filename host target path failover
state type or hardware path adapter port priority
-----
up secondary /dev/dsk/c6t0d1 fcd0 1a 0 (Wrong)
up primary /dev/dsk/c9t0d1 fcd1 1a 1
up primary /dev/dsk/c4t0d1 fcd0 1b 2
up secondary /dev/dsk/c11t0d1 fcd1 1b 3
```

4. Reconfigure the PV-Links path order by entering the following command:

```
# ontap_config_paths
```

5. Enter the `ontap_config_paths` command for the second time.

You must run this command again because it did not set the correct path order when you entered it the first time.

```
# ontap_config_paths
```

6. Verify that the paths are configured in the correct order by entering the following command:

```
# 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.

```
# sanlun lun show -p
ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
LUN: 1
LUN Size: 100m
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-38
Mode: 7
VG: /dev/test_vg
Multipath Policy: A/P
Multipath Provider: Native
-----
host controller /dev/dsk controller PVlinks
path path filename host target path failover
state type or hardware path adapter port priority
-----
up primary /dev/dsk/c9t0d1 fcd1 1a 0
up primary /dev/dsk/c4t0d1 fcd0 1b 1
up secondary /dev/dsk/c11t0d1 fcd1 1b 2
up secondary /dev/dsk/c6t0d1 fcd0 1a 3
```

Commands you can use to check for problems

You can use several commands to check your system and look for problems.

The following commands are especially useful in tracking down problems:

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

The HP-UX `ioscan` command

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

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

- 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 ONTAP LUN dev/dsk/c31t0d5 /dev/rdisk/c31t0d5
disk 3 0/4/0/0.1.6.0.26.0.5 sdisk    NO_HW DEVICE   ONTAP LUN /dev/dsk/c34t0d5 /dev/rdisk/c34t0d5
disk 2 0/6/2/0.1.13.0.26.0.5 sdisk    CLAIMED DEVICE ONTAP LUN /dev/dsk/c30t0d5 /dev/rdisk/c30t0d5
disk 4 0/6/2/0.1.14.0.26.0.5 sdisk    NO_HW DEVICE   ONTAP LUN /dev/dsk/c35t0d5 /dev/rdisk/c35t0d5
```

The `sanlun` utility

You can use the `sanlun` utility to check the status of the HBA, LUN, and PV-Link path; to determine the `/dev/dsk/filename` association with the storage system LUN; and to 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 storage system is primary or secondary.
- Report missing paths from the volume group. (You can use the `ontap_config_paths` utility to easily correct any missing paths.)
- Report the PV-Links path ordering.
- Get information about a takeover.

You can use `sanlun` with the `-p` option to check your paths. 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.

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

```
# sanlun lun show -p

ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
LUN: 1
LUN Size: 100m
Controller CF State: Cluster Enabled
Controller Partner: f3070-210-38
Mode: 7
VG: /dev/test_vg
Multipath Policy: A/P
Multipath Provider: Native
-----
host controller /dev/dsk
path path filename host controller PVlinks
state type or hardware path adapter target path failover
priority
-----
up secondary /dev/dsk/c6t0d1 fcd0 1a 0 (Wrong)
up secondary /dev/dsk/c11t0d1 fcd1 1b 1 (Wrong)
up primary /dev/dsk/c9t0d1 fcd1 1a 2
up primary /dev/dsk/c4t0d1 fcd0 1b not in VG
```

Note: You can repair the ordering of the paths by running the `ontap_config_paths` utility.

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 information:

- 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

ONTAP Path: f3070-210-37:/vol/bootvol/lun-10
LUN: 1
LUN Size: 100m
Controller CF State: Partner Takeover
Controller Partner: f3070-210-37
Mode: 7
VG: /dev/test_vg
Multipath Policy: A/P
Multipath Provider: Native
-----
host controller /dev/dsk
path path filename host controller PVlinks
state type or hardware path adapter target path failover
priority
-----
down /dev/dsk/c9t0d1 0
down /dev/dsk/c4t0d1 1
```

up	secondary	/dev/dsk/c11t0d1	fcd1	1b	2
up	secondary	/dev/dsk/c6t0d1	fcd0	1a	3

The enable_ontap_pvlinks script

You can use the `enable_ontap_pvlinks` script to change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing other vendors' LUNs.

Note: Before running the `enable_ontap_pvlinks` script, ensure that Data ONTAP LUNs are visible on the host.

The `enable_ontap_pvlinks` command has the following format:

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

To set the multipathing policy to active/passive for all Data ONTAP LUNs without disturbing other vendors' LUNs enter the following command:

```
# enable_ontap_pvlinks set
```

To display the current multipathing policy for all Data ONTAP LUNs, enter the following command:

```
# enable_ontap_pvlinks show
```

Replacing a LUN on HP-UX 11iv3

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

Steps

1. Stop all the I/O on the LUN device that must 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 you unmapped the LUN in the previous step.
5. Enter `ioscan` command on the host.

During the `ioscan` operation, you will see the following message in the host `syslog` file

```
(/var/adm/syslog/syslog.log):
```

```
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.
```

6. Locate the paths that turned to NO_HW because of this replacement by entering the following command:

```
# ioscan -fnNC lun_path|grep NO_HW
```

7. Manually validate the change in the WWID of the LUN by entering the following command on the paths:

```
# scsimgr replace_wwid -H lun_path
```

Example

```
# 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
```

Note: The `scsimgr` man page provides more information for this command.

Copyright and trademark information

Copyright ©1994 - 2012 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2012 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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

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

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

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Trademark information

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service

Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

- A**
- ALUA
 - disabling Native MPIO ALUA 55
 - disabling Native MPIO ALUA for Veritas 16
 - supported with HP-UX 16
 - APM
 - available from Symantec 50
 - failover path 51
 - I/O error handling 51
 - I/O error processing 51
 - obtaining from Symantec 52
 - path failures 51
 - upgrading 54
 - used with Veritas 50
 - ASL
 - upgrading 54
 - ALUA 55
 - available from Symantec 50
 - error messages 51
 - obtaining from Symantec 52
 - used with Veritas 50
 - using VxVM 58
- C**
- configurations
 - supported for Host Utilities 14
- D**
- depot 20
 - determining the ASL version 54
- E**
- enable_ontap_pvlinks
 - installed by Host Utilities 10
 - Environment 11
 - environments
 - finding information on different Host Utilities environments 12
 - error messages
 - ASL 51
- F**
- FCoE
 - converged network adapters 13
 - traditional FC 13
- H**
- HBA
 - displaying information with sanlun 46
 - Host Utilities
 - contents 10
 - instructions for the different environments 12
 - planning installation 18
 - supported configurations 14
 - HP-UX 11iv2
 - getting I/O policy information 57
 - HP-UX 11iv3
 - getting I/O policy information 57
 - HP-UX Host Utilities 10
 - HP-UX LVM
 - PVLinks 16
- I**
- I/O policy
 - finding for Veritas 57
 - igroup
 - configuration overview 27
 - uses WWPN 13
 - installation
 - overview 18
 - planning 18
 - instructions
 - finding information on different environments 12
 - iSCSI protocol
 - implementing with Host Utilities 13
 - instructions for that Host Utilities environment 12
- L**
- LUNs
 - adding in Veritas environment 38
 - configuration overview 27
 - displaying with sanlun 43
 - LVM
 - instructions for that Host Utilities environment 12
- M**
- man pages

- installed by Host Utilities 10
- MPIO
 - instructions for that Host Utilities environment 12
- multipathing
 - using sanlun with Veritas 58
- Multipathing 15

N

- Native MPIO ALUA
 - disabling 55
- nSANity
 - diagnostic utility 10
- ntap_config_paths
 - installed by Host Utilities 10

P

- Protocol 11
- protocols
 - Fibre Channel (FC) 12
 - iSCSI 12

R

- restore policy
 - recommended value 57

S

- SAN boot, FC
 - SAN boot LUNs 61
- SAN booting
 - advantages 17
- san_version
 - installed by Host Utilities 10
- sanlun utility
 - displaying array types 58
 - displaying LUNs 43
 - displaying multipathing for Veritas 58
 - installed by Host Utilities 10
- space reclamation
 - command for reclaiming 59
 - viewing 59
- storage 28
- storage system
 - direct-attached 12
 - Ethernet network 12

- fabric-attached 12
- Symantec
 - getting the ASL and APM 52
 - provides ASL, APM 50

T

- thin provisioning
 - ASL required 59
 - viewing 59

U

- Using FC, iSCSI, and FCoE 11
- Using HP-UX 11

V

- Veritas
 - restore daemon settings 57
 - adding a LUN 38
 - APM 50
 - ASL 50
 - disabling Native MPIO ALUA 55
 - displaying multipathing 58
 - I/O policy 57
 - thin provisioning 59
 - using ALUA with HP-UX 16
 - Volume Manager 16
- Veritas DMP
 - instructions for that Host Utilities environment 12
 - SAN booting instructions 17
- Veritas Storage Foundation
 - Array Policy Module, APM 52
 - Array Support Library, ASL 52
- vgcreate 35
- VxVM
 - Veritas Volume Manager 16
 - displaying LUN paths 58
 - VxVM, command 52

W

- WWID 74
- WWPN
 - getting with sanlun fcp show adapter command 13
 - required for igroups 13



NA 210-05569_A0, Printed in USA

GA32-0764-03

