

IBM System Storage N series



# HP-UX Host Utilities 5.2 Installation and Setup Guide



# Contents

|   |           |
|---|-----------|
| <b>Preface .....</b>  | <b>7</b>  |
| Supported features .....  | 7         |
| Websites .....  | 7         |
| Getting information, help, and service .....                                    | 7         |
| Before you call .....   | 8         |
| Using the documentation .....   | 8         |
| Hardware service and support .....  | 8         |
| Firmware updates .....  | 9         |
| How to send your comments .....   | 9         |
| <b>Changes to this document: December 2011 .....</b>                            | <b>11</b> |
| <b>The HP-UX Host Utilities .....</b>   | <b>13</b> |
| Supported HP-UX environments and protocols .....                                | 14        |
| How to find instructions for your HP-UX environment .....                       | 15        |
| Protocols and configurations supported by Host Utilities .....                  | 15        |
| The FC protocol .....   | 16        |
| The FCoE protocol .....   | 16        |
| The iSCSI protocol .....  | 17        |
| Supported configurations .....  | 17        |
| Features supported by the Host Utilities .....                                  | 18        |
| Multipathing and the HP-UX Host Utilities .....                                 | 18        |
| HBAs and the HP-UX Host Utilities .....   | 19        |
| Volume managers on the HP-UX Host Utilities .....                               | 19        |
| HP-UX configurations that support ALUA .....                                    | 19        |
| SAN booting and the Host Utilities .....  | 20        |
| <b>Quick start checklist for planning the Host Utilities installation .....</b> | <b>23</b> |
| <b>Installing the Host Utilities software .....</b>                             | <b>27</b> |
| Removing the HP-UX Host Utilities software .....                                | 28        |
| <b>(iSCSI) Installing and Configuring the iSCSI Software Initiator .....</b>    | <b>31</b> |
| (iSCSI) Downloading the iSCSI Software Initiator .....                          | 31        |
| (iSCSI) Installing the iSCSI Software Initiator .....                           | 31        |
| (iSCSI) Verifying the iSCSI Software Initiator installation .....               | 32        |
| (iSCSI) Configuring the iSCSI Software Initiator .....                          | 32        |

|   |           |
|---|-----------|
| Recording the host iSCSI initiator node name .....                                    | 33        |
| (iSCSI) Configuring the discovery targets .....                                       | 33        |
| <b>LUN configuration .....</b>  | <b>35</b> |
| Overview of LUN configuration and management .....                                    | 35        |
| Discovering LUNs on an HP-UX host .....   | 36        |
| (HP-UX 11iv2) Configuring LUNs for use with LVM .....                                 | 39        |
| (HP-UX 11iv3) Configuring LUNs for use with LVM .....                                 | 43        |
| <b>The sanlun utility .....</b>   | <b>47</b> |
| Displaying host LUN information with sanlun .....                                     | 47        |
| Displaying host HBA information with sanlun .....                                     | 50        |
| <b>(Veritas DMP) Array Support Library and Array Policy Module .....</b>              | <b>53</b> |
| (Veritas DMP) What the Array Support Library and the Array Policy Module<br>are ..... | 53        |
| (Veritas DMP) What an ASL array type is .....   | 54        |
| (Veritas DMP) Information provided by the ASL .....                                   | 54        |
| (Veritas DMP) How to upgrade the ASL and APM .....                                    | 54        |
| (Veritas DMP) ASL and APM installation overview .....                                 | 55        |
| (Veritas DMP) Determining the ASL version .....                                       | 55        |
| (Veritas DMP) How to get the ASL and APM .....  | 56        |
| (Veritas DMP) Tasks to perform before you uninstall the ASL and APM ...               | 56        |
| (Veritas DMP) Installing the ASL and APM software .....                               | 57        |
| (Veritas DMP, HP-UX 11iV3) Thin provisioning and space reclamation .....              | 58        |
| (Veritas DMP, HP-UX 11iV3) Disabling Native MPIO ALUA .....                           | 59        |
| (Veritas DMP) Using VxVM to display available paths .....                             | 61        |
| (Veritas DMP) Adding a new LUN to VxVM environment .....                              | 61        |
| (Veritas DMP) The Veritas DMP restore daemon requirements .....                       | 64        |
| (Veritas DMP) The Veritas DMP I/O policy .....  | 65        |
| (Veritas DMP) How sanlun displays the array type .....                                | 65        |
| (Veritas DMP) Using sanlun to obtain multipathing information .....                   | 66        |
| (Veritas DMP) Information about ASL error messages .....                              | 66        |
| <b>SAN boot for LVM and VxVM in FC environments .....</b>                             | <b>69</b> |
| (LVM) Prerequisites for creating a SAN boot using LVM .....                           | 70        |
| (LVM) Creating a bootable device on an Itanium system .....                           | 70        |
| (LVM) Determining the size of the boot LUN .....                                      | 71        |
| (LVM) Creating and mapping a LUN to an igroup .....                                   | 72        |
| (LVM) Discovering the boot LUN as a host device .....                                 | 73        |

|  |            |
|--|------------|
| (LVM) Migrating an operating system disk using MirrorDisk/UX on PA-RISC systems .....                        | 74         |
| (LVM) Mirroring the boot disk to the boot LUN .....  | 74         |
| (LVM) Setting the host's primary and secondary boot paths .....  | 77         |
| (LVM) Installing the operating system using an Ignite Server or installation CD-ROM on PA-RISC systems ..... | 79         |
| (LVM) Identifying paths to the boot LUN .....  | 80         |
| (LVM) Installing the operating system on the boot LUN .....  | 81         |
| (LVM) Testing the secondary boot path for PA-RISC systems .....  | 82         |
| (VxVM) Creating a Veritas-controlled SAN Boot LUN on HP-UX .....   | 83         |
| (VxVM) Creating the boot LUN using a cold install .....  | 83         |
| <b>Troubleshooting .....</b>   | <b>85</b>  |
| Host setup checklist .....   | 85         |
| Storage system setup checklist .....   | 86         |
| Connectivity problems when using LVM .....   | 86         |
| Migrating a configuration from non-ALUA to ALUA without rebooting the host .....                             | 87         |
| Situations affecting volume groups and disk devices .....  | 88         |
| Application I/O failures on Veritas VxVM .....   | 88         |
| (Veritas DMP) Enclosure-based naming is not working correctly .....  | 88         |
| Problems with volume group multipathing .....  | 88         |
| (LVM) Troubleshooting sanboot configurations with PA-RISC systems .....                                      | 89         |
| Handling a takeover and giveback .....   | 90         |
| Events that occur if a host reboots during a takeover .....  | 91         |
| Recovering from a host reboot during a takeover .....  | 91         |
| Enable the host to identify more than eight LUNs per target .....  | 93         |
| Commands you can use to check for problems .....   | 96         |
| The HP-UX ioscan command .....   | 96         |
| The sanlun utility .....   | 96         |
| The enable_ontap_pvlinks script .....  | 98         |
| Replacing a LUN on HP-UX 11iv3 .....   | 98         |
| <b>Copyright information .....</b>   | <b>101</b> |
| <b>Trademark information .....</b>   | <b>103</b> |
| <b>Index .....</b>   | <b>105</b> |



# 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 N series support Web site (accessed and navigated as described in [Websites](#) on page 7).

## 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 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/](http://www.ibm.com/storage/nas/)
- The IBM System Storage N series support Web site 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 Web site, refer to the following publicly accessible web page:  
[www.ibm.com/storage/support/nseries/](http://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](http://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](http://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 Web site (accessed and navigated as described in [Websites](#) on page 7) 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 Web site (accessed and navigated as described in [Websites](#) on page 7).

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 Web site.

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](http://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/](http://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 Web site (accessed and navigated as described in *Websites* on page 7).

**Note:** If you do not see new firmware updates on the N series support Web site, 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 e-mail to [starpubs@us.ibm.com](mailto: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: December 2011

---

Several changes have been made to this document since it was published for the HP-UX Host Utilities 5.1. The 5.2 version of the HP-UX Host Utilities adds a few essential supports, and this document is updated to include that information.

### December 2011 update

In December 2011, this document has been updated to add the information regarding the following supports:

- FCoE support with HP-UX 11iv3.  
This support contains the FCoC driver for use with Fibre Channel Over Ethernet HBAs.
- Support for more than 500 SAN LUNs.  
The earlier versions of SAN LUN displayed only up to 500 LUNs even if you mapped more than 500 LUNs on an HP-UX host. This limitation is now overcome.

### December 2010 update

Several changes have been made to this document since it was published for the HP-UX Host Utilities 5.0. At that time, the document was updated to include following information:

- Information about the diagnostic utilities has been removed. Previously, the Host Utilities included the `controller_info`, `filer_info`, `hpux_info`, and `switch_info` utilities. Starting with the HP-UX Host Utilities 5.1, the diagnostic utilities are no longer included with the Host Utilities. They have been replaced by the nSANity Diagnostic and Configuration Data Collector program. The nSANity program is not part of the Host Utilities. If needed, please contact N series support for assistance to perform diagnostics.
- Information about using Veritas Storage Foundation 5.0.1 RP2 with HP-UX 11iv3 has been added.  
Only ALUA configurations are supported with HP-UX 11iv3 and Veritas Storage Foundation. However, you must disable Native MPIO ALUA on Veritas LUNs and use the ALUA supplied with the Symantec Array Support Library (ASL).
- Information about using thin provisioning and storage reclamation with Veritas Storage Foundation 5.0.1 RP2 and HP-UX 11iv3. You can use the `vxdisk` and `fsadm` commands to perform space reclamation on a LUN, disk group, enclosure, or mounted Veritas file system.

### Related tasks

*(Veritas DMP, HP-UX 11iv3) Disabling Native MPIO ALUA* on page 59

### Related references

*(Veritas DMP, HP-UX 11iV3) Thin provisioning and space reclamation* on page 58



## 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` utility, which lets you enable the active/passive multipathing policy for Physical Volume Links (PV-Links) with Data ONTAP LUNs.
- The `ontap_config_paths` utility, which lets you add missing paths to LVM volume groups and place them in the correct order when you are using PV-Links.
- The `sanlun` utility, which helps you to manage LUNs and HBAs.
- 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. If needed, please contact N series support for assistance to perform diagnostics.

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:

- This 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*

## 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 online N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7).

The following table summarizes key aspects of the main environments:

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

**Related information**

*IBM N series Support page - [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)*

**How to find instructions for your HP-UX environment**

Many of the instructions in this manual 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:

| <b>Qualifier</b> | <b>The section that follows applies to</b>   |
|------------------|--|
| (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 DMP)    | 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.

**Protocols and configurations supported by Host Utilities**

The Host Utilities provide support for Fibre Channel, 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.

For more information about using the protocols with your storage system, see the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.

For more details on supported topologies, including diagrams, see the *Fibre Channel and iSCSI Configuration Guide* for your version of Data ONTAP.

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

### Related information

[N series Interoperability Matrices Web site - www.ibm.com/systems/storage/network/interophome.html](http://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 *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.

### Related tasks

[Displaying host HBA information with `sanlun`](#) on page 50

### Related references

[Overview of LUN configuration and management](#) on page 35

## 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.

In general, you can configure and use FCoE connections the same way you use traditional FC connections.



## 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 *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.

### Related tasks

[Recording the host iSCSI initiator node name](#) on page 33

### Related references

[Overview of LUN configuration and management](#) on page 35

## Supported configurations

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

The *Fibre Channel and iSCSI Configuration Guide* provides detailed information, including diagrams, about the supported topologies. There is also configuration information in the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP. Refer to those documents for complete information about configurations and topologies.

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 an active/active 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 are not a limitation in Ethernet iSCSI topologies. Refer to the Ethernet switch vendor documentation for specific recommendations and best practices.

## Features supported by the Host Utilities

The Host Utilities support many features and configurations available with HP-UX hosts and storage systems running Data ONTAP. Your specific environment affects what the Host Utilities support.

Some of supported features include:

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

The N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7) describes which features are supported for which configurations.

### Related information

*[N series Interoperability Matrices Web site - www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.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 go away.

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 Host Utilities support a number of HBAs.

Make sure the supported HBAs are installed before you install the Host Utilities. Normally, the HBAs should have the correct firmware and FC driver. To determine the firmware and FC driver for your system, run the appropriate administration tool for your HBA.

**Note:** For details on the specific HBAs that are supported and the required firmware and FC drives, see the N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7).

### Related information

*N series Interoperability Matrices Web site - [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)*

## Volume managers on 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<br><b>Note:</b> 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 N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7) Interoperability Matrix.

### Related tasks

*(HP-UX 11iv3) Configuring LUNs for use with LVM* on page 43

*Discovering LUNs on an HP-UX host* on page 36

*Migrating a configuration from non-ALUA to ALUA without rebooting the host* on page 87

### Related references

*The `enable_ontap_pvlinks` script* on page 98

### Related information

*N series Interoperability Matrices Web site - [www.ibm.com/systems/storage/network/interophome.html](http://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 in both the Veritas DMP and PV-Links/Native MPIO 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.

The steps you must perform to create a SAN boot LUN differ based on your Host Utilities environment. If you decide to set up SAN booting, ensure that you use the instructions for your environment.



# Quick start 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.

Regardless of your experience level, planning how you will install and configure the Host Utilities is a good practice.

The detailed steps for each of the tasks presented in the checklist are provided later in this guide.

**Note:** Occasionally there are known problems that can affect your system setup. Please review 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 discovered after this guide was produced.

## Task 1: Make sure the prerequisites for installing and setting up the Host Utilities have been met

1. Verify that your system setup is correct. Check the N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7) for the most current information about system requirements. This includes:
  - Host operating system version and appropriate updates.
  - HBAs and drivers, or software initiators, model and version.
  - **(Veritas)** Veritas Storage Foundation.  
Set up the Veritas Volume Manager (VxVM) and install the Array Support Library (ASL) and Array Policy Module (APM), if these modules weren't installed with your version of Veritas Storage Foundation. 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. You must also set the Veritas DMP restore daemon interval to 60.
    - Note:** If you need to manually install the ASL and APM, make sure you first have the Veritas VxVM installed. You can get the ASL and APM from the Symantec Web site.
2. Verify that your storage system has Data ONTAP installed and is:
  - Running the correct, licensed protocol for your environment.
  - Set up to work with the host and the HBAs or software initiators, as needed by your protocol.
3. **(FC/FCoE)** If you are using a switch, verify that it is:
  - Set up correctly.

- Zoned.
- Cabled correctly according to the instructions in the *Fibre Channel and iSCSI Configuration Guide* 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, see the *Fibre Channel and iSCSI Configuration Guide* for your version of Data ONTAP.

4. Confirm that the host and the storage system can communicate.

### **Task 2: Install the Host Utilities**

1. Download a copy of the compressed file containing the Host Utilities from the N series support Web site at [www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/).
2. If you are upgrading the Host Utilities, remove any currently installed version of the Host Utilities.
3. Use the `gunzip` command to uncompress the file.
4. Install the Host Utilities software using the `swinstall` command.

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

5. Log out and log back in to enable the updates that the installation script makes to the system search paths.

### **(iSCSI) Task 3: Configure the iSCSI protocol**

1. Download the iSCSI software initiator from the HP Web site.
2. Use the `swlist` command to verify that the file 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, follow the instructions provided by the multipathing vendor to set it up to work with iSCSI.

4. Record the host's iSCSI node name.
5. Add information to the targets file.
6. Set up target discovery.
7. **(Optional)** Set up Challenge Handshake Authentication Protocol (CHAP) on the host and the storage system. To set up CHAP, follow the instructions provided by 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, set it up to work with the igroups.

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 Host Utilities software

---

The Host Utilities software is provided as a standard HP-UX depot file. After your host is set up, install the SAN Toolkit software package. You can install it the same way you normally install an HP-UX depot file.

## Before you begin

Keep the following in mind before you install the Host Utilities software:

- You must verify that your host system meets the Host Utilities requirements.
- If you are upgrading the HP-UX Host Utilities 5.2 from an earlier version, you must first uninstall the earlier version.
- **(iSCSI only)** If you are using iSCSI, you must install your iSCSI software initiator before you install the Host Utilities.

## Steps

1. Download the HP-UX Host Utilities file `ibm_hpux_host_utilities_5.2_ia_pa.depot.gz` to your HP-UX host.
2. Uncompress the `ibm_hpux_host_utilities_5.2_ia_pa.depot.gz` file by entering the following command:

```
# gunzip ibm_hpux_host_utilities_5.2_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 /tmp/hpux/ibm_hpux_host_utilities_5.2_ia_pa.depot
Ontap_santoolkit
```

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 `/tmp/hpux` directory on an HP-UX host called `machine1`.

```
===== 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**

[\(iSCSI\) Downloading the iSCSI Software Initiator](#) on page 31

[\(iSCSI\) Installing the iSCSI Software Initiator](#) on page 31

[Removing the HP-UX Host Utilities software](#) on page 28

## 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 5.2 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.

**Example**

```
# 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 `machine1`:

**Example**

Enter the following command to get the output as illustrated by the example:

```
# swremove Ontap_santoolkit
```

```
===== 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_titanium,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) Downloading the iSCSI Software Initiator

You can download the iSCSI Software Initiator from an HP Web site.

### Steps

1. Go to the HP Web site.
2. In the Search box, enter iSCSI Software Initiator.
3. Click **Receive for Free** in the search result.
4. Select the iSCSI Software Initiator version that you want to download.

**Note:** More information about installing the iSCSI Software Initiator is available on the HP Web site.

### Related information

[The HP Web site](#)

## (iSCSI) Installing the iSCSI Software Initiator

After you download the iSCSI Software initiator, you must verify that the file 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 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 downloaded 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* provide 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 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 several executables to the root path and run the `iscsiutil` command.

### Steps

1. Add the `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.

## 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 `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.

2. 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

1. 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.

*tcp-port* is the TCP port component of the discovery target network portal. The default iSCSI TCP port number is 3260

*portal-grp-tag* is the target portal group tag. The default target portal group tag for discovery targets is 1.

2. View the configured discovery targets by entering the following command:

```
iscsiutil -p -D
```

3. Discover the operational target devices by entering the following command:

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

4. 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 which 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                            | Discussion  |
|---------------------------------|---|
| Create and map 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:</p> <ul style="list-style-type: none"> <li>• <b>(FC/FCoE)</b> The WWPN for the HBA ports on the host. You can use the <code>sanlun fcp show adapter</code> command to get the WWPN.</li> <li>• <b>(iSCSI)</b> The iSCSI initiator node name for the host. You can use the <code>iscsiutil -l</code> command to get this.</li> </ul> <p>After you create the igroup, you must create LUNs on the storage system and map the LUNs to the igroup.</p> <p>To create and map igroups and LUNs, use the <code>lun setup</code> command or enter a series of individual commands such as <code>lun create</code>, <code>igroup create</code>, and <code>lun map</code>.</p> <p>The <i>Data ONTAP Block Access Management Guide for iSCSI and FC</i> for your version of Data ONTAP provides complete information in creating igroups and LUNs.</p> |
| Enable ALUA                     | <p>If your environment supports ALUA, you must set it up to work with igroups. To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.</p>   |

| Task                                 | Discussion  |
|--------------------------------------|---|
| Discover LUNs                        | After you create the LUN and map it to your igroup, you must discover the LUN as a host device.   |
| Configure 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 50

[Recording the host iSCSI initiator node name](#) on page 33

### Related information

[N series Interoperability Matrices Web site - www.ibm.com/systems/storage/network/interphome.html](http://www.ibm.com/systems/storage/network/interphome.html)

## Discovering LUNs on an HP-UX host

To configure and manage the LUNs, the LUNs must be discovered by the host.

### Before you begin

- Reboot the host. The host automatically discovers new LUNs when you reboot the host.
- If rebooting the host is not a reasonable action for the storage system, perform the steps listed here. These steps enable the host to discover LUNs without a host reboot.

**Note:** The HP-UX Host Utilities provide Veritas support with HP-UX 11iv2 and HP-UX 11iv3. The iSCSI Software Initiator is not supported with Veritas Storage Foundation.

### Steps

1. Log in as root on the host.
2. To get the host 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 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

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

- 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.

```
# ioscan -fnC disk
Class I H/W Path Driver          S/W State H/W Type Description
-----
disk 9 0/1/1/0/1/1.2.0          sdisk CLAIMED DEVICE HP
36.4GATLAS10K3_36_SCA
      /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk 1 0/4/2/0/4/0.1.8.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN
disk 2 0/4/2/0/4/0.1.8.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN
disk 3 0/4/2/0/4/0.1.12.0.0.0.2 sdisk CLAIMED DEVICE NETAPP LUN
disk 4 0/4/2/0/4/0.1.12.0.0.0.4 sdisk CLAIMED DEVICE NETAPP LUN
```

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

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

```

# sanlun lun show -p all
f3040-210-142:/vol/testvol/lun50 (LUN 99)
Agile Filename: /dev/rdisk/disk501
    4g (4294967296)    lun state: GOOD
Controller_CF_State: Cluster Enabled  Controller Partner: f3040-210-143
Multipath_Provider: None
-----
host      controller
path      path          /dev/dsk host  primary  partner
state     type          filename HBA      controller controller
-----
up        primary      /dev/dsk/c2t12d3 fcd2    0c
up        secondary   /dev/dsk/c4t12d3 fcd2          0c
up        secondary   /dev/dsk/c9t12d3 fcd3          0d
up        primary      /dev/dsk/c7t12d3 fcd3    0d

```

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

5. **(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.

6. **(HP-UX 11iv3 February 2007 release)** Display 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 once 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 by entering the following command:

```
mknod
```

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 0x01000.

```
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:

```
vgdisplay -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 /dev/vg_ntap01
--- Volume groups ---
VG   Name /dev/vg_ntap01
VG   Write Access read/write
VG   Status available
Max  LV 255
Cur LV 0
Open LV 0
Max  PV 16
Cur PV 1
Act  PV 1
Max  PE per PV 1016
VGDA 2
```

```

PE Size (Mbytes) 4
Total PE 511
Alloc PE 0
Free PE 511
Total PVG 0
Total Spare PVs 0
Total Spare PVs in use 0
#
# lvcreate -l 511 /dev/vg_ntap01
Logical volume "/dev/vg_ntap01/lvol1" has been successfully created with
character device "/dev/vg_ntap01/rlvol1".
Logical volume "/dev/vg_ntap01/lvol1" has been successfully extended.
Volume Group configuration for /dev/vg_ntap01 has been saved in
/etc/lvmconf/vg_ntap01.conf
#

```

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 `vgdisplay -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. The output shows the multipathing policy and the PV-Links order of the paths to the LUN.

```
# sanlun lun show -p
f3170-210-97:/vol/lun_hpux_20_h2/lun2 (LUN 5) VG: /dev/vg01
  2g (2147483648)      lun state: GOOD
Controller_CF_State: Cluster Enabled  Controller Partner: f3170-210-96
Multipath_Policy: A/P
Multipath_Provider: Native
-----
host   controller      /dev/dsk      primary      partner      PVlinks
path   path              filename host   controller controller path failover
state  type              or Hardware Path HBA   port        port        priority
-----
up     primary          /dev/dsk/c11t0d5 fcd2 0c
up     primary          /dev/dsk/c9t0d5 fcd1 0d
up     secondary        /dev/dsk/c13t0d5 fcd2
up     secondary        /dev/dsk/c7t0d5 fcd1
                                0d
                                0c
```

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

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

**Before you begin**

Before mapping the LUNs, enable ALUA on the igroup by entering the following command:

```
igroup set igroup_name alua on
```

**Steps**

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

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

You must perform this process once for each LUN when you have multiple paths to that LUN.

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

```
mkdir /dev/directory_name
```

**Example**

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

3. Create a volume group with one primary path 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 sets up a primary path to the volume group.

```
# 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:

```
vgdisplay -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:

```
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 /dev/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          1.0
VG Max Size         1036320m
VG Max Extents      259080
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/rlvol1" has been successfully created with
character device "/dev/vg_ntap01/rlvol1".
Logical volume "/dev/vg_ntap01/rlvol1" 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 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
#

```

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. The output shows the multipathing policy and the PV-Links order of the paths to the LUN.

```
# sanlun lun show -p
f3170-210-97:/vol/lun_hpux_12_h2/lun2 (LUN 103)
Agile Filename: /dev/rdisk/disk520 VG: /dev/n_vg
14.0g (15033434112) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3170-210-96
Multipath_Policy: A/A
Multipath_Provider:Native
-----
host controller /dev/dsk primary partner PVlinks
path path filename host controller controller path failover
state type or Hardware Path HBA port port port priority
-----
up primary /dev/dsk/c7t12d7 fcd0 0c 0
up primary /dev/dsk/c13t12d7 fcd1 0d 0
up secondary /dev/dsk/c9t12d7 fcd0 0d 1
up secondary /dev/dsk/c11t12d7 fcd1 0c 1
```

# 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 `sanlun` 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 file_name | all | controller_name | controller_name:pathname]
```

```
sanlun lun show -p [-v] [ all | controller_name | controller_name:pathname]
```

```
sanlun lun show -b [ all | controller_name | controller_name:pathname]
```

**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` provides information about the primary and secondary paths available to the LUN when you are using multipathing. It lists the multipathing software used to access the LUN (HP-UX PV-Links, HP-UX Native MPIO, or Veritas DMP) and the multipathing policy. For example, if you are using PV-Links, the `-p` option provides information about path priority.

If you are using DMP, the `-p` option provides information about the multipathing policy.

If you are using HP-UX Native MPIO, the `-p` option provides information about the native multipathing policy.

`all` displays information for all LUNs that are currently attached to the host.

`-b` prints a brief listing of the LUNs.

`-d` specifies the special device file on the host (*host\_device\_filename*). For HP-UX 11iv2, the file would be `-d /dev/rdisk/c4t0d7`; for HP-UX 11iv3, it would be `/dev/rdisk/disk56`.

*controller\_name* is the name of the target storage system.

*controller\_pathname* 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 policy associated with the LUN by entering the following command:

```
sanlun lun show -p
```

If you are using the HP-UX LVM, this command is useful if you need to view path ordering for PV-Links or troubleshoot a problem with path ordering.

- Display the summary listing of the storage system LUNs associated with the host device `/dev/rdisk/c3t4d2` by entering the following command:

```
sanlun lun show -d /dev/rdisk/c3t4d2
```

- 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 list of paths between the host and a specific LUN (for example, `lun0`) on the storage system `controllerA` by entering the following command:

```
sanlun lun show controllerA:lun_path
```

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 example, `lun0`) `controllerA` by entering the following command:

```
sanlun lun show controllerA:lun_path
```

### Example

The examples that follow show the type of output you would see on a system using LVM and on a system using DMP as well as when you use the `sanlun lun show` command with the `-v` option. You should keep the following information in mind as you look at the LVM and DMP output.

- A primary path type indicates a connection that goes to the storage system that owns the LUN. The path can be either through a SAN switch or a direct-attached connection. This is the preferred route for I/O to the LUN. A secondary path type indicates a path through a



secondary storage system (a storage system that does not own the LUN). The primary path type is used only when the primary path to the LUN is unavailable.

- The controller port column refers to the port on the storage system on which the LUN is accessed. It always represents a primary path. The partner port column refers to the port on the partner in the cluster. It always represents a secondary path.

### 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 before running the `ontap_config_paths` utility. In this example, LUNs are managed by LVM and PV-Links. The PV-Links failover priority column shows the PV-Links order of the paths to the LUN.

```
# sanlun lun show -p
f3170-210-97:/vol/lun_hpux_20_h2/lun2 (LUN 5) VG: /dev/vg01
  2g (2147483648) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3170-210-96
Multipath_Policy: A/P
Multipath_Provider: Native
-----
host controller /dev/dsk primary partner PVlinks
path path filename host controller controller path failover
state type or Hardware Path HBA port port port priority
-----
up primary /dev/dsk/c11t0d5 fcd2 0c 0
up primary /dev/dsk/c9t0d5 fcd1 0d 1
up secondary /dev/dsk/c13t0d5 fcd2 0d 2
up secondary /dev/dsk/c7t0d5 fcd1 0c 3
```

### 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
f3040-210-142:/vol/LUN_NDU_VRTS_/lun40 (LUN 89)
Agile Filename: /dev/rdisk/disk505 DMP-NODE: (N52000_94)
  4g (4294967296) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3040-210-143
Multipath_Provider: Veritas
-----
host controller /dev/dsk primary partner
path path filename HBA controller controller
state type or Hardware Path HBA port port
-----
up primary /dev/dsk/c6t11d1 fcd5 0d
up secondary /dev/dsk/c8t11d1 fcd5 0d
up primary /dev/dsk/c10t11d1 fcd6 0c
up secondary /dev/dsk/c12t11d1 fcd6 0c
```

### Example using LVM with ALUA while legacy mode is disabled

On an HP-UX 11iv3 September 2007 release or later system using LVM with ALUA enabled on the controller igroup and legacy I/O nodes and legacy DSFs disabled on the host, `sanlun` returns the following output:

```
# sanlun lun show -p

f3040-210-143:/vol/mvol93/lun93 (LUN 0)
Agile Filename: /dev/rdisk/disk802      VG: /dev/vg01
    1.6g (1702887424)      lun state: GOOD
Controller_CF_State: Cluster Enabled   Controller Partner: f3040-210-142
Multipath_Policy: A/A
Multipath_Provider: Native
-----
host    controller /dev/dsk          primary   partner   HP A/A
path    path          filename         host     controller controller path
failover
state   type          or Hardware Path HBA      port      port      priority
-----
up      primary       0/4/1/0/4/0.0x500a098197297dd9.0x4000000000000000
                                fcd7     0c          0
up      primary       0/5/1/0/4/0.0x500a098297297dd9.0x4000000000000000
                                fcd0     0d          0
up      secondary    0/4/1/0/4/0.0x500a098287297dd9.0x4000000000000000
                                fcd7          0d          1
up      secondary    0/5/1/0/4/0.0x500a098187297dd9.0x4000000000000000
                                fcd0          0c          1
```

### Example using `-v`

When you enter the `sanlun` command with the `-v` option, you might see the following output:

```
# sanlun lun show -v
controller:          lun-pathname          device filename  adapter
protocol            lun size             lun state
f3040-210-143:    /vol/LUN_NDU_NATIVE_/lun1  /dev/rdisk/c7t0d0  fcd3
FCP                4g (4294967296)      GOOD
Serial number: Hn/NVo/JzSfq
Controller FCP nodename:500a098087097de0  Controller FCP
portname:500a098287097de0
Controller adapter name: v.0d
Controller IP address: 10.72.210.143
Controller volume name:LUN_NDU_NATIVE_    FSID:0x2ee6668
Controller qtree name:/vol/LUN_NDU_NATIVE_ ID:0x0
Controller snapshot name: ID:0x0
```

## Displaying host HBA information with `sanlun`

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

### Steps

1. Log in as root on the host.

2. At the host command line, display HBA information by entering the following command:

**sanlun**

This command has the following format:

```
sanlun fcp show adapter [ -c | [ -v ] [adapter_name | all ]]
```

-c displays configuration information that you can use to create igroups.

-v produces verbose output.

all lists information for all FC adapters.

*adapter\_name* lists adapter information for the adapter you specify.

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

Using the `# sanlun fcp show adapter` with the option `adapter` produces the following output.

```
# sanlun fcp show adapter
td0 WWPN:50060b000023e2ec
td1 WWPN:50060b000023e2cc
```

Adding the `-c` option, so that you enter `# sanlun fcp show adapter -c` produces the following output.

```
# sanlun fcp show adapter -c
igroup create -f -t hpux "Node-2" 50060b000023e2ec 50060b000023e2cc
```

Adding the `-v` option, so that you enter `# sanlun fcp show adapter -v` produces the following output.

```
# sanlun fcp show adapter -v
adaptername:      td0
WWPN:             50060b000023e2ec
WNNN:            50060b000023e2ed
drivername:      td
model:           A6795A
modeldescription: FibreChannelMassStorageAdapter (PCI) -2Gig
serialnumber:    NotAvailable
hardwareversion: 2.3
driverversion:   @(#) libtd.aHPFibreChannelTachyonXL2DriverB.
11.23.0512
$Date:           2005/09/2012:22:47$Revision:r11.23/1
firmwareversion: NotAvailable
Numberofports:   1
porttype:        Fabric
portstate:       Operational
supportedspeed:  2GBit/sec
negotiatedspeed: 2GBit/sec
OSdevicename:    /dev/td0
adaptername:     td1
WWPN:            50060b000023e2cc
WNNN:            50060b000023e2cd
drivername:      td
model:           A6795A
```

```

modeldescription:FibreChannelMassStorageAdapter(PCI)-2Gig
serialnumber:      NotAvailable
hardwareversion:   2.3
driverversion:     @(#)libtd.aHPFibreChannelTachyonXL2DriverB.11.23.0512
$Date:
2005/09/2012:22:47$Revision:r11.23/1
firmwareversion:  NotAvailable
Numberofports:    1
porttype:         Fabric
portstate:        Operational
supportedspeed:   2GBit/sec
negotiatedspeed:  2GBit/sec
OSdevicename:    /dev/td1

```

FCoE adapter produces the following output when using the command # `sanlun fcp show adapter -v`.

```

#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

```

## (Veritas DMP) Array Support Library and Array Policy Module

---

If you are using the Veritas Volume Manager with Dynamic Multipathing (DMP) for multipathing support on an HP-UX system, ensure that the VxVM software, Veritas patch bundle, and Array Support Library (ASL) and Array Policy Module (APM) are installed. Veritas Software Foundation -11iv2 supports version 5.0 and 11iv3 supports version 5.0.1 and 5.1 SP1 onwards.

### (Veritas DMP) What the Array Support Library and the Array Policy Module are

The ASL and APM for N series storage systems are provided by Symantec. You must download the ASL and APM from the Symantec Web site.

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

To determine which versions of the ASL and APM you need for your version of the host operating system, check the appropriate interoperability matrix for your N series product. This information is updated frequently. When you know which version you need, go to the Symantec Web site and download the ASL and APM.

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

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

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

#### Related information

*[N series Interoperability Matrices Web site -www.ibm.com/systems/storage/network/interphome.html](http://www.ibm.com/systems/storage/network/interphome.html)*

## (Veritas DMP) What an ASL array type is

The ASL reports information about the multipathing configuration to the DD and specifies the configuration as a disk array type.

The configuration is identified as one of the following disk array types:

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

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

## (Veritas DMP) Information provided by the ASL

The ASL provides enclosure-based naming information and array information about SAN-attached storage systems.

The ASL lets you obtain the following information about the LUNs:

- Enclosure name.  
With enclosure-based naming, the name of the Veritas disk contains the model name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to VxVM about SAN-attached storage systems, instead of referring to them as Just a Bunch of Disks (JBOD) devices or raw devices. The enclosure-based naming feature used by VxVM creates a disk name based on the name of its enclosure, or disk array, and not a raw device name.

For details about system management, see *Veritas Volume Manager Administrator's Guide*. Veritas documents are available at Veritas Storage Foundation DocCentral, which, at the time this document was prepared, was available online at <http://sfdoccentral.symantec.com/>.

## (Veritas DMP) How to upgrade 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.

## (Veritas DMP) ASL and APM installation overview

If you are using DMP with Veritas Storage Foundation 5.0 or later, you must install the ASL and the APM.

The basic installation of the ASL and the APM involves the following tasks:

- Verify that your configuration meets system requirements. See the N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7) for current information about the system requirements.
- If you currently have the ASL installed, determine its version.
- If you need to install a newer version of the ASL and APM, remove the older versions before you install the new versions.

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

- Obtain the new ASL and the APM.
- Follow the instructions in the Symantec TechNote to install the new version of the ASL and APM.

### Related information

*N series Interoperability Matrices Web site - [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)*

## (Veritas DMP) Determining the ASL version

If you currently have the ASL installed, you should check its version to determine whether you need to update it.

### Step

1. 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
=====
libvxautoraid.sl       vm-5.0-rev-1          5.0
libvxcscovrts.sl       vm-5.0-rev-1          5.0
libvxCLARiON.sl        vm-5.0-rev-1          5.0
libvxemc.sl            vm-5.0-rev-2          5.0
libvxxfc60.sl          vm-5.0-rev-1          5.0
libvxva.sl             vm-5.0-rev-1          5.0
libvxhdsalua.sl        vm-5.0-rev-1          5.0
libvxhpalua.sl         vm-5.0-rev-1          5.0
libvxshark.sl          vm-5.0-rev-1          5.0
libvxhds9980.sl        vm-5.0-rev-1          5.0
libvxhdsusp.sl         vm-5.0-rev-2          5.0
```

|                                |                           |     |
|--------------------------------|---------------------------|-----|
| <code>libvxibm8k.sl</code>     | <code>vm-5.0-rev-1</code> | 5.0 |
| <code>libvxxp1281024.sl</code> | <code>vm-5.0-rev-1</code> | 5.0 |
| <code>libvxxp12k.sl</code>     | <code>vm-5.0-rev-2</code> | 5.0 |
| <code>libvxnetapp.sl</code>    | <code>vm-5.0-rev-3</code> | 5.0 |

## (Veritas DMP) How to get the ASL and APM

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

To get the ASL and APM, go to the Symantec Web site and download them.

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 TechNote is provided on the N series Interoperability Matrices Web site at [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html).

## (Veritas DMP) Tasks to perform before you uninstall 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.

## (Veritas DMP) Example of uninstalling the ASL and the APM

The following is an example of uninstalling the ASL and the APM when you have Veritas Storage Foundation 5.0.

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.

```
# swremove VRTSNTAPapm
===== 05/20/08 18:28:17 IST BEGIN swremove SESSION
(non-interactive) (jobid=hpx_19-0149)
* Session started for user "root@hpx_19".
* Beginning Selection
* Target connection succeeded for "hpx_19:/".
* Software selections:
VRTSNTAPapm.APM_FILES,l=/,r=5.0,v=VERITAS,fr=5.0,fa=HPUX_
B.11.23_PA
* Selection succeeded.
* Beginning Analysis
* Session selections have been saved in the file
```



```

"/.sw/sessions/swremove.last".
* The analysis phase succeeded for "hpux_19:/".
* Analysis succeeded.
* Beginning Execution
* The execution phase succeeded for "hpux_19:/".
* Execution succeeded.
NOTE: More information may be found in the agent
logfile using the
command "swjob -a log hpux_19-0149 @ hpux_19:".
===== 05/20/08 18:28:35 IST END swremove SESSION
(non-interactive)
(jobid=hpux_19-0149)
# swremove VRTSNTAPas1
===== 05/20/08 18:29:01 IST BEGIN swremove SESSION
(non-interactive) (jobid=hpux_19-0150)
* Session started for user "root@hpux_19".
* Beginning Selection
* Target connection succeeded for "hpux_19:/".
* Software selections:
VRTSNTAPas1.ASL_FILES,l=/,r=5.0,a=HPUX_
B.11.23_IA/PA,v=VERITAS,fr=5.0,fa=HP-UX_B.11.23_PA
* Selection succeeded.

```

## (Veritas DMP) Installing the ASL and APM software

To install a fresh version of the ASL and APM that you downloaded from Symantec involves several steps.

### Before you begin

- Make sure you obtain the ASL and APM TechNote, which you can view at the Symantec Web site. The TechNote contains the Symantec instructions for installing the ASL and APM.
- You should have your LUNs set up before you install the ASL and APM.

### Steps

1. Log in to the VxVM system as the root user.
2. If you have your N series storage configured as JBOD in your VxVM configuration, remove the JBOD support for the storage by entering:

```
vxddladm rmjbod vid=NETAPP
```

3. Verify that you have downloaded the correct version of the ASL and APM by checking the N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7). If you do not already have the correct version or the ASL and APM TechNote, you can follow the link in the matrix to the correct location on the Symantec Web site.
4. Install the ASL and APM according to the installation instructions provided by the ASL/APM TechNote on the Symantec Web site.
5. If your host is connected to N series storage, verify your installation by entering:

```
vxddmpadm listenclosure all
```

By locating the Enclosure Type in the output of this command, you can verify the installation. The output shows the model name of the storage device if you are using enclosure-based naming with VxVM.

6. If your host is not connected to storage, use the following command:

```
vxddladm listsupport all
```

7. Verify that the APM is installed by entering following command:

```
vxddmpadm listapm all
```

### After you finish

After you install the ASL and APM, you should perform the following procedures:

- If you have Data ONTAP 7.1 or later, it is recommended that you change the cfmode setting of your clustered systems to single-image mode, and then reconfigure your host to discover the new paths to the disk.
- On the storage system, create LUNs and map them to the igroups containing the WWPNs of the host HBAs.
- On the host, discover the new LUNs and configure them to be managed by VxVM.

### Related information

*N series Interoperability Matrices Web site - [www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.html)*

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

When you are running HP-UX 11iV3 and Veritas Storage Foundation 5.0.1 RP2, you can use the VxFS space reclamation feature. For more information about configurations that support this feature, see N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7).

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

Contact Symantec Support to get this ASL.

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

- DMP node
- enclosure
- Veritas 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 mach30400_12
Reclaiming thin storage on:
Disk mach30400_12 : Done
```

This command reclaims space at the DMP enclosure level.

```
vxdisk reclaim mach30400
Reclaiming thin storage on:
Disk mach30400_12 : Done!
Disk mach30400_13 : Done!
Disk mach30400_14 : Done!
Disk mach30400_15 : Done
```

This command reclaims space at the Veritas disk group level.

```
vxdisk reclaim dg1
Reclaiming thin storage on:
Disk mach30400_15 : Done!
Disk mach30400_12 : Done!
Disk mach30400_13 : Done!
Disk mach30400_14 : Done
```

This command reclaims space at the file system level.

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

## (Veritas DMP, 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 get_attr -D /dev/rdisk/disk460 -a alua_enabled
```

### Example

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

```
# 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 -p get_attr all_lun -a device_file -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 19

## Related information

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

## (Veritas DMP) Using VxVM to display available paths

If a LUN is being managed by VxVM, then you can use VxVM to display information about available paths to that 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.

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 `vxdisk list` command.

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

```
vxddmpadm getsubpaths ctrl=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.

## (Veritas DMP) Adding a new LUN to VxVM environment

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 an HP-UX host*.
2. Enable VxVM to detect the new LUN by entering the following command:

```
vxddctl enable
```

3. 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
controller_0 auto:none - - online invalid
controller_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 (for example, `c0t0d3`) is displayed.

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

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

*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

```
# vxdisk list
DEVICE TYPE DISK GROUP STATUS
Disk_0 auto:LVM - - LVM
Disk_1 auto:LVM - - LVM
controller_0 auto:none - - online invalid
controller_1 auto:none - - online invalid
# vxdisksetup -i FAS30200_0 format=cdsdisk
# vxdisksetup -i FAS30200_1 format=cdsdisk
# vxdisk list
DEVICE TYPE DISK GROUP STATUS
Disk_0 auto:LVM - - LVM
Disk_1 auto:LVM - - LVM
controller_0 auto:cdsdisk - - online
controller_1 auto:cdsdisk - - online
```

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

```
vxvg 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.

6. Create disks or add disks to the disk group by entering the following command:

```
vxvg -g dg_name adddisk diskname=device_name
```

#### Example

```
# vxvg -g dg1 adddisk dg1-disk2=controller_1
# vxdisk list
DEVICE TYPE DISK GROUP STATUS
Disk_0 auto:LVM - - LVM
Disk_1 auto:LVM - - LVM
```

```
controller_0 auto:cdsdisk dg1-disk1 dg1 online
controller_1 auto:cdsdisk dg1-disk2 dg1 online
```

7. 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.

8. 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.

9. 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
```

11. 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-UX11iv2.

**Note:** The `sanlun` utility only displays the native multipathing policy. To see the multipathing policy for other vendors, you must use vendor-specific commands.

### HP-UX 11iv3 output

```
# sanlun lun show -p
controller1:/vol/TP_LUN_h2_/lun2 (LUN 11)
Agile Filename: /dev/rdisk/disk561      DMP-NODE:
      5g (5368709120)      lun state: GOOD
Controller_CF_State: Cluster Enabled  Controller Partner: f3040-210-143
Multipath_Provider: Veritas
-----
host      controller
path      path          /dev/dsk host  primary  partner
state     type          filename HBA      port     controller
-----
up        primary      /dev/dsk/c7t1d3 fcd1    0d
up        secondary   /dev/dsk/c9t1d3 fcd1           0d
```

```
up      primary      /dev/dsk/c3t1d3 fcd2  0c
up      secondary   /dev/dsk/c5t1d3 fcd2                0c
```

**HP-UX 11iv3 with FCoE LUN.** The command `# sanlun lun show -p` displays the following output for the FCoE LUN.

```
/vol/FCOE_HPUX_/lun1 (LUN 0)
Agile Filename: /dev/rdisk/disk559      DMP-NODE:
 10g (10737418240)   lun state: GOOD
Controller_CF_State: Cluster Enabled  Controller Partner:
Multipath_Provider: Veritas
-----
host      controller
path      path                /dev/dsk host   primary  partner
state     type                filename HBA         port    controller
-----
up        primary          /dev/dsk/c37t0d0 fcoc4   0d
up        primary          /dev/dsk/c43t0d0 fcoc5   0c
up        secondary       /dev/dsk/c39t0d0 fcoc4
up        secondary       /dev/dsk/c41t0d0 fcoc5                0c
                                     0d
```

### HP-UX 11iv2 output

```
# sanlun lun show -p
controller2:/vol/verts_tst_/lun2 (LUN 1) DMP-NODE: (F30400_1)
 5g (5368709120)   lun state: GOOD
Controller_CF_State: Cluster Enabled  Controller Partner: f3040-210-142
Multipath_Provider: Veritas
-----
host      controller
path      path                /dev/dsk host   primary  partner
state     type                filename HBA         port    controller
-----
up        primary          /dev/dsk/c21t0d1 fcd2   0d
up        secondary       /dev/dsk/c19t0d1 fcd2                0d
up        primary          /dev/dsk/c17t0d1 fcd1   0c
up        secondary       /dev/dsk/c15t0d1 fcd1                0c
```

## (Veritas DMP) 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)

It is a good practice to check the *Release Notes* to see if these values have changed since this document was produced.



## (Veritas DMP) 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 `vxdmpadm getattr enclosure <storage_system> iopolicy`. This command displays output similar to the following:

```
# /usr/sbin/vxdmpadm getattr enclosure fas30400 iopolicy
ENCLR_NAME      DEFAULT          CURRENT
=====
bas30400        MinimumQ        MinimumQ
```

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

```
bash-2.05# /usr/sbin/vxdmpadm list dmpnode dmpnodename=disk558
dmpdev          = SS30400_9
state           = enabled
enclosure       = SS30400
cab-sno         = 30008136
asl             = libvxnetapp.sl
vid             = NETAPP
pid             = LUN
array-name      = SS3040
array-type      = ALUA
iopolicy        = MinimumQ
avid            = -
lun-sno         = Hn/NVoZqlQNO
udid            = NETAPP%5FLUN%5F30008136%5FHn%2FNVoZqlQNO
dev-attr        = tprclm
###path         = name state type transport ctlr hwpath aportID aportWWN attr
path            = c7t1d0 enabled secondary FC c7 0/4/1/0.0x500a098287097de0 50:0a:
09:82:87:09:7d:e0 - -
path            = c3t1d0 enabled secondary FC c3 0/4/1/1.0x500a098187097de0 50:0a:
09:81:87:09:7d:e0 - -
path            = c9t1d0 enabled(a) primary FC c9 0/4/1/0.0x500a098297097de0 50:0a:
09:82:97:09:7d:e0 - -
path            = c5t1d0 enabled(a) primary FC c5 0/4/1/1.0x500a098197097de0 50:0a:
09:81:97:09:7d:e0 - -
```

## (Veritas DMP) How sanlun displays the array type

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

The Host Utilities installation process places the `sanlun` utility in the `/opt/Ontap/santools/bin/` directory.

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, then 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 DMP) Using sanlun to obtain multipathing information

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 the ASL is installed and the LUN is controlled by VxVM, the output of the `sanlun` command displays the `Multipath_Policy` as either.

If the LUN is not controlled by a volume manager, then the `Multipath_Policy` is none and the `Multipathing-provider` is none.

### 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 DMP) 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 <code>vxdisk</code> list, but it is not usable. |

| <b>Message severity</b> | <b>Definition</b>  |
|-------------------------|--|
| 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. |



# SAN boot for LVM and VxVM in FC environments

You can set up a SAN boot to work in an LVM and Veritas VxVM environment using the FC protocol.

**Note:** At the time this document was produced, LVM and Veritas VxVM environments only supported SAN boot with the FC protocol. The Veritas LVM and Veritas VxVM environments did not support SAN boot with the iSCSI protocol.

The procedures in this chapter assume the following:

- You are configuring a PA-RISC server.

**Note:** Itanium systems use partitions for the boot device, so some of the procedures might be different for Itanium systems.

- The host is booting from an existing internal operating system disk. It is not a diskless server.
- You manage the boot LUN with HP-UX LVM or Veritas VxVM.
- The SAN Toolkit software is installed on the host.

To configure the HP-UX operating system on the boot LUN using LVM, you must perform these tasks:

1. Migrate an existing operating system disk to the boot LUN by using MirrorDisk/UX.
2. Install the operating system by using an Ignite Server or the HP installation CD-ROM.
3. Install the operating system on the boot LUN from the internal root disk by using the `drd clone` command.

To configure the HP-UX operating system on the boot LUN using Veritas, you must perform these tasks:

1. Create a Veritas-controlled SAN boot on HP-UX.
2. Perform a cold install directly from the HP-UX Core operating system DVD.

The sections that follow provide steps for configuring a SAN boot and installing a new operating system image on it. You can apply these steps to most configurations.

Other methods exist for creating a SAN boot. This guide does not describe other methods for creating bootable LUNs, such as creating configurations that boot multiple volumes or use diskless servers.

**Note:** Solutions and components are qualified continually. To verify that SAN booting is supported in your configuration, see the N series Service and Support Web site.

## Related information

*N series Interoperability Matrices Web site - [www.ibm.com/systems/storage/network/interphome.html](http://www.ibm.com/systems/storage/network/interphome.html)*

## (LVM) Prerequisites for creating a SAN boot using LVM

You must set up your HP-UX host system and install the Host Utilities before you create a SAN boot LUN.

- If you are installing the operating system on the boot LUN, create a LUN that is large enough to provide enough space for the version of HP-UX you are using. For details, see the documentation for your version of HP-UX.
- If you are mirroring an existing host boot disk, create a LUN that is at least as large as the allocated space used by the existing system boot disk.

## (LVM) Creating a bootable device on an Itanium system

Itanium systems use three partitions for the boot device.

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

The following `sanlun` sample output from an Itanium system using LVM and PV-Links shows the paths to a bootable LUN.

```
controller1:/vol/vol1/bootlun (LUN 101) VG: /dev/vg00
 18g (19327352832) lun state: GOOD
Filer_CF_State: Cluster Enabled Filer Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
host filer          primary partner PVlinks
path path          /dev/dsk      host filer filer path
failover
state type          filename      HBA  port    port    priority
-----
up    primary         /dev/dsk/c7t12d5s2 td0  4a     0
up    primary         /dev/dsk/c13t12d5s2 td1  5a     1
```

|    |           |                     |     |    |   |
|----|-----------|---------------------|-----|----|---|
| up | secondary | /dev/dsk/c11t12d5s2 | td1 | 5b | 2 |
| up | secondary | /dev/dsk/c5t12d5s2  | td0 | 4b | 3 |

## (LVM) Determining the size of the boot LUN

You must determine the size needed for the boot LUN before you create the bootable LUN.

### Steps

1. On the host, obtain information about the root volume group (vg00) and determine which disk is used as the system disk by entering the following command:

```
vgdisplay -v vg00
```

#### Example

The following example shows the path to the system disk:

```
--- Physical volumes ---
PV Name      /dev/dsk/c1t15d0
PV Status    available
Total PE     4340
Free PE      2136
Autoswitch   On
```

The output listed under physical volumes shows the paths to the system disk.

2. On the host, display the size of the system disk by entering the following command:

```
diskinfo /dev/rdisk/raw_device_node
```

*raw\_device\_node* is the system disk device listed under the physical volumes in the output of the `vgdisplay` command you used in Step 1.

#### Example

```
diskinfo /dev/rdisk/c1t15d0
SCSI describe of /dev/rdisk/c1t15d0:
 vendor: HP 18.2G
  product id: MAN3184MC
   type: direct access
  size: 17783240 Kbytes
 bytes per sector: 512
```

The output shows that the size of the system disk is approximately 18 GB. Your boot LUN must be at least 18 GB.

## (LVM) Creating and mapping a LUN to an igroup

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

### Before you begin

If you do not have an igroup set up, you must create one. The *Data ONTAP Block Access Management Guide for iSCSI and FC* provides detailed information about creating igroups and LUNs.

### Steps

1. On the storage system command line, enter the following command:

```
lun create -s size -t hpux lun_path
```

*size* is the size of the LUN, which you determined when you ran the `diskinfo` command. You can use one of the following characters to establish the size: c (bytes), w (words or double bytes), b (512), k (kilobytes), m (megabytes), g (gigabytes), or t (terabytes).

*lun\_path* is the LUN's path name, which includes the volume and qtree.

### Example

```
lun create -s 18g -t hpux /vol/bootvol/bootlun1
```

2. Map the LUN to the igroup by entering the following command:

```
lun map lun_path igroup_name LUN_ID
```

*lun\_path* is the path name of the LUN you created.

*igroup\_name* is the name of the igroup you created.

*LUN\_ID* is the identification number that you specify. The host associates the LUN ID with the location and path name of the LUN. If you do not specify a LUN ID, Data ONTAP automatically assigns one. The default LUN ID begins with 0 and increments by 1 for each additional LUN that is created. If you have existing LUNs, Data ONTAP assigns the next available LUN ID.

### Example

```
lun map /vol/bootvol/bootlun1 boothost 0
```

3. Verify that the LUN is online and mapped to the igroup by entering the following command:

```
lun show -v
```

### Example

```
mac_1> lun show -v
/vol/bootvol/bootlun1 18g (19327352832) (r/w,online,
mapped)
```



```
Serial#: OdCtSYv7Y6w3
Share: none
Space Reservation: enabled
Multiprotocol Type: hpux
Maps: boothost=0
```

## (LVM) Discovering the boot LUN as a host device

After you create the boot LUN and map it to the igroup, the host must discover the LUN as a host device.

### Before you begin

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

### Steps

1. Log in as root on the host.
2. Check whether the host has discovered the newly created LUNs by entering the following command:

```
ioscan -fnC disk
```

### Example

The output from an `ioscan` command provides the following information:

- No device special file exists for LUNs yet.
- The LUN is visible to the host and the S/W State of each path to the LUN is CLAIMED, which means that the path is available.

```
# ioscan -fn -C disk
Class I H/W Path Driver          S/W State H/W Type Description
=====
disk 9 0/1/1/0/1/1.2.0          sdisk CLAIMED DEVICE HP
36.4GATLAS10K3_36_SCA
      /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk 1 0/4/2/0/4/0.1.8.0.0.0.2 sdisk CLAIMED DEVICE NETAPP
LUN
disk 2 0/4/2/0/4/0.1.8.0.0.0.4 sdisk CLAIMED DEVICE NETAPP
LUN
disk 3 0/4/2/0/4/0.1.12.0.0.0.2 sdisk CLAIMED DEVICE NETAPP
LUN
disk 4 0/4/2/0/4/0.1.12.0.0.0.4 sdisk CLAIMED DEVICE NETAPP
LUN
```

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, use the `insf -e` command to install the device special files.

4. Determine which of the following methods you will use to configure the HP-UX operating system on the boot LUN:
  - Migrate an existing system boot disk to the LUN by using MirrorDisk/UX.
  - Reinstall the operating system using an Ignite Server or an installation CD-ROM.

## (LVM) Migrating an operating system disk using MirrorDisk/UX on PA-RISC systems

To migrate an existing system boot disk to the operating system disk using MirrorDisk/UX on PA-RISC systems, you must mirror the existing boot disk to the new boot LUN and set the host's primary and secondary boot paths.

### (LVM) Mirroring the boot disk to the boot LUN

You must mirror the existing disk to migrate the system disk to the newly created LUN.

#### Steps

1. On the host, determine which paths to the boot LUN are primary paths by entering the following command:

```
sanlun lun show -p
```

The `sanlun` command is in the `/opt/Ontap/santools/bin` directory. Use a primary path from one HBA to the LUN to create a physical volume and add it to the root volume group.

#### Example

In this example, the device `c24t0d0` is a primary path you can use for the root volume group. This device is used throughout this procedure example as the boot LUN.

```
# sanlun lun show -p
boothost:/vol/vol1/bootlun1 (LUN 0)
18g (19327352832) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3040-210-143
Multipath_Provider: None
-----
host controller          primary partner
path path    /dev/dsk      host controller controller
state type    filename      HBA port      port
-----
up  secondary /dev/dsk/c44t0d0 td1          5b
up  primary   /dev/dsk/c46t0d0 td1 5a
up  secondary /dev/dsk/c26t0d0 td0          4b
up  primary   /dev/dsk/c24t0d0 td0 4a
```

2. **(HP-UX 11iv3 February 2007)** The HP-UX 11iv3 February 2007 release does not support ALUA, so the default multipathing policy for all the disk storage is active/active. If you are using the HP-UX 11iv3 February 2007 release, you might need to change the multipathing policy to active/passive for Data ONTAP LUNs. You can do this by using the `enable_ontap_pvlinks` utility. Enter the following command:

```
enable_ontap_pvlinks set
```

The `enable_ontap_pvlinks` utility lets you change the multipathing policy to active/passive for Data ONTAP LUNs without disturbing other vendors' LUNs.

**Note:** If you are using HP-UX 11iv2, you do not need to use this utility.

3. **(HP-UX 11iv3 February 2007)** Display the current multipathing policy for all Data ONTAP LUNs by entering the following command:

```
enable_ontap_pvlinks show
```

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

```
pvcreate -fB /dev/rdisk/disk_device
```

`disk_device` is the raw disk device on the primary path to the LUN. Use the raw device node for the device you found in Step 1.

`-fB` forces the creation of the physical volume and creates space for a boot area on the LUN.

```
pvcreate -fB /dev/rdisk/c24t0d0
```

### Example

The following example uses the `pvcreate` command with the device `c24t0d0`.

```
pvcreate -fB /dev/rdisk/c24t0d0
```

5. Add the LUN to the root volume group (`vg00`). Use a block disk device node, not a raw disk device. Enter the following command:

```
vgextend vg00 /dev/dsk/disk_device
```

`disk_device` is the disk device file for the primary path to the LUN. Use the device you found in Step 1.

```
vgextend vg00 /dev/dsk/c24t0d0
```

### Example

The following example uses the `vgextend` command to add the LUN to the root volume group (`vg00`)

```
vgextend vg00 /dev/dsk/c24t0d0
```

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

**Example**

The following example uses the utility to add the rest of the paths to the LUN to the root volume group (vg00):

```
ontap_config_paths vg00
```

7. Copy the boot utilities to the boot LUN by entering the following command:

```
mkboot /dev/rdisk/disk_device
```

*disk\_device* is the device used for the boot LUN. For example, c24t0d0.

8. Ensure that the host will boot if either the boot LUN or the system disk is unavailable by entering the following commands:

```
mkboot -a "hpux -lq" /dev/dsk/system_disk_device
```

```
mkboot -a "hpux -lq" /dev/dsk/ntap_device
```

*system\_disk\_device* is the path to the system disk device node. For example, c1t15d0.

*ntap\_device* is the device node on the primary path to the boot LUN. For example, c24t0d0.

9. Display the logical volumes on the existing system boot disk (vg00) by entering the following command:

```
vgdisplay -v vg00
```

The logical volumes are displayed under the “Logical volumes” line in the command output.

**Example**

The following partial output from the `vgdisplay` command shows how two logical volumes are listed:

```
--- Logical volumes ---
LV Name      /dev/vg00/lvol1
LV Status    available/syncd
LV Size      (Mbytes) 300
Current LE   75
Allocated PE 75
Used PV      1
LV Name      /dev/vg00/lvol2
LV Status    available/syncd
LV Size (Mbytes) 2048
Current LE   512
Allocated PE 512
Used PV      1
```

10. For each logical volume, mirror the logical volume from the existing system boot disk (vg00) to the boot LUN by entering the following command:

```
lvextend -m 1 lv_path /dev/dsk/ntap_device
```

`-m 1` specifies the number of mirror copies.

*lv\_path* is the path to the logical volume. For example, /dev/vg00/lvol2.

`ntap_device` is the device on the primary path to the boot LUN. For example, `c24t0d0`.

### Example

This example mirrors logical volume 1 to the boot LUN (`/dev/dsk/c24t0d0`).

```
lvextend -m 1 /dev/vg00/lvol1 /dev/dsk/c24t0d0
The newly allocated mirrors are now being synchronized. This operation will
take some time. Please wait...
Logical volume "/dev/vg00/lvol1" has been successfully extended. Volume Group
configuration for /dev/vg00 has been saved in /etc/lvmconf/vg00.conf
```

11. Update the Boot Device Reserved Area (BDRA) by entering the following command:

```
lvlnboot -R
```

12. Using `vi` or another text editor, add a line to the `/stand/bootconf` file to identify the boot disk to the SDUX utility. Use the path to the block device node.

```
1 /dev/dsk/ntap_device
```

`ntap_device` is the device used for the boot LUN.

```
1 /dev/dsk/c24t0d0
```

## (LVM) Setting the host's primary and secondary boot paths

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

### About this task

After you mirror the boot disk to the boot LUN, you must set the primary and secondary boot paths for the host by using the `setboot` command on the host. When you use this command, you specify hardware paths to a device. Before you can run the `setboot` command, you must identify the primary and secondary paths to the boot LUN by using the `sanlun` command, and then identify which hardware paths they correspond to by using the `ioscan` command.

When you define the primary boot path, you must ensure that you use the device that you identified. This is the device you added to the root volume group and to which you copied the HP-UX utilities.

### Steps

1. Display the paths to the boot LUN by entering the following command:

```
sanlun lun show -p
```

### Example

```
# sanlun lun show -p
boothost:/vol/vol1/bootlun (LUN 0) VG: /dev/vg00
18g (19327352832) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3170-210-96
Multipath_Policy: A/P
Multipath_Provider: Native
-----
host controller primary partner
```

| path state | path type | /dev/dsk filename | host HBA | controller port | controller port |
|------------|-----------|-------------------|----------|-----------------|-----------------|
| up         | primary   | /dev/dsk/c24t0d0  | td0      | 4a              | 0               |
| up         | primary   | /dev/dsk/c46t0d0  | td1      | 5a              | 1               |
| up         | secondary | /dev/dsk/c44t0d0  | td1      | 5b              | 2               |
| up         | secondary | /dev/dsk/c26t0d0  | td0      | 4b              | 3               |

In this example from a host using LVM, the primary path to the boot LUN from the host is `/dev/dsk/c24t0d0`. This is the path to which you mirrored logical volumes and copied HP-UX utilities.

A secondary path to the boot LUN from the host is `/dev/dsk/c44t0d0`. These paths are used for the remainder of this procedure example.

- Determine which hardware paths correspond to the primary and secondary paths to the boot LUN that you displayed in Step 1 by entering the following command:

```
ioscan -funC disk
```

#### Example

```
# ioscan -funC disk
Class I H/W Path Driver S/W State H/W Type Description
=====
==
disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC
/dev/dsk/c1t15d0 /dev/rdisk/c1t15d0
disk 14 0/4/0/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c24t0d0 /dev/rdisk/c24t0d0
disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c26t0d0 /dev/rdisk/c26t0d0
disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c46t0d0 /dev/rdisk/c46t0d0
disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c44t0d0 /dev/rdisk/c44t0d0
```

In this example, the hardware path for the primary path to the boot LUN is `0/4/0/0.1.4.0.0.0.0`. Use this hardware path for the host's primary boot path.

The hardware path for the secondary path to the boot LUN is `0/6/2/0.1.7.0.0.0.0`. Use this hardware path for the host's secondary boot path.

- Set the host's primary boot path by entering the following command:

```
setboot -p prim_hw_path_to_lun
```

`prim_hw_path_to_lun` is the primary hardware path to the boot LUN from the host's HBA1, displayed in Step 2.

#### Example

The following example shows how to set the primary boot path to the boot LUN by using the hardware path that corresponds to `/dev/dsk/c24t0d0`, the primary path.

```
setboot -p 0/4/0/0.1.4.0.0.0.0
```

4. Set the host's secondary boot path by entering the following command:

```
setboot -a sec_hw_path_to_lun
```

*sec\_hw\_path\_to\_lun* is the secondary hardware path to the boot LUN from the host's HBA2, displayed in Step 2.

#### Example

The following example shows how to set the alternate boot path to the boot LUN by using the hardware path that corresponds to `/dev/dsk/c44t0d0`, the secondary path.

```
setboot -a 0/6/2/0.1.7.0.0.0.0
```

5. Test the primary boot path by shutting down and rebooting the system:

```
shutdown -r -y 0
```

**Note:** If automatic booting is disabled (set to OFF) for your system, you must boot the system from the Boot Control Handler (BCH) menu.

6. Perform the following task:

| If...                         | Follow these instructions...   |
|-------------------------------|--|
| You cannot reboot the system. | Troubleshooting sanboot configurations with PA-RISC systems (LVM only) |
| You can reboot the system.    | Testing the secondary boot path for PA-RISC systems (LVM only)         |

## (LVM) Installing the operating system using an Ignite Server or installation CD-ROM on PA-RISC systems

You use the same procedure for installing the operating system using an Ignite Server and or an installation CD-ROM on PA-RISC systems.

To install the operating system on the boot LUN, you must specify a primary path to the boot LUN as the root disk. The installation process uses the `setboot` command to define the primary path to the boot LUN as the host's primary boot path.

After the installation is complete, you manually enter the `setboot` command to define a secondary boot path. Therefore, you must identify the primary and secondary paths to the boot LUN, and their corresponding hardware paths, before you begin the installation.

## (LVM) Identifying paths to the boot LUN

To install the operating system on the boot LUN, you must identify the paths to the boot LUN.

### Steps

1. Display the primary and secondary paths to the boot LUN by entering the `sanlun` command, which is in the `/opt/Ontap/santools/bin` directory.

```
sanlun lun show -p
```

#### Example

In this example, a primary path to the boot LUN from the host is `/dev/dsk/c24t0d0`. A secondary path to the boot LUN from the host is `/dev/dsk/c44t0d0`. These paths are used as the primary and secondary paths for the remainder of this procedure example.

```
# sanlun lun show -p
boothost:/vol/vol1/bootlun (LUN 0)
18g (19327352832) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: f3170-210-96
Multipath_Policy: A/P
Multipath_Provider: Native
```

| host  | controller | primary          | partner    |    |
|-------|------------|------------------|------------|----|
| path  | path       | host             | controller |    |
| state | type       | HBA              | controller |    |
|       |            | port             | port       |    |
| up    | primary    | /dev/dsk/c24t0d0 | td0        | 4a |
| up    | primary    | /dev/dsk/c46t0d0 | td1        | 5a |
| up    | secondary  | /dev/dsk/c44t0d0 | td1        | 5b |
| up    | secondary  | /dev/dsk/c26t0d0 | td0        | 4b |

2. Determine which hardware paths correspond to the primary and secondary paths to the boot LUN that you displayed in Step 1 by entering the following command:

```
ioscan -func disk
```

#### Example

In this example, the hardware path for the primary path to the boot LUN is `0/4/0/0.1.4.0.0.0.0`. Use this hardware path for the host's primary boot path.

The hardware path for the secondary path to the boot LUN is `0/6/2/0.1.7.0.0.0.0`. Use this hardware path for the host's secondary boot path.

```
# ioscan -func disk
Class I H/W Path Driver S/W State H/W Type Description
=====
==
disk 0 0/0/1/1.15.0 sdisk CLAIMED DEVICE HP 18.2GMAN3184MC
/dev/dsk/clt15d0 /dev/rdisk/clt15d0
disk 14 0/4/0/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c24t0d0 /dev/rdisk/c24t0d0
disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c26t0d0 /dev/rdisk/c26t0d0
```



```
disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c46t0d0 /dev/rdisk/c46t0d0
disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c44t0d0 /dev/rdisk/c44t0d0
```

- Record the primary and secondary hardware device paths.

During installation, you use a primary path to the LUN to define the root disk (vg00). After you install the operating system, you use a secondary path to define the host's secondary boot path.

## (LVM) Installing the operating system on the boot LUN

To install the operating system on the boot LUN, you must have identified the primary path to the boot LUN as the root disk.

### Steps

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

| If...   | Then...   |
|---|---|
| <b>You are installing from a CD-ROM</b>         | Follow the HP instructions to boot from the CD-ROM.   |
| <b>You are installing from an Ignite Server</b> | <ul style="list-style-type: none"> <li>Reboot the host by entering the following command:<br/><b>shutdown -r -y 0</b></li> </ul> <p><b>Note:</b> If automatic booting is enabled (set to ON), you must interrupt the boot process to display the BCH menu. If automatic booting is disabled (set to OFF), the system comes up and displays the BCH menu.</p> <ul style="list-style-type: none"> <li>When the system is up again, enter the following command at the BCH Main Menu prompt:<br/><b>bo lan.server_IP</b><br/><i>server_IP</i> is the IP address of the Ignite Server.</li> </ul> <pre>Main Menu: Enter command or menu &gt; bo lan.10.105.36.44 Interact with IPL (Y, N, or Cancel)?&gt; n</pre> |

- In the HP-UX installation, configure system parameters, such as network services, file-system characteristics, and other configuration details.

Ensure that you define the root disk as the primary path to the boot LUN. In the Ignite-UX interface, you define the root disk on the Basic tab.

When you define the primary path to the boot LUN as the root disk, the installation script complete the following steps:

- Configures the root volume group (vg00) and adds the remaining path to the LUN to the root volume group.

- Uses the `setboot -p` command to define the primary path to the boot LUN as the host's primary boot path.

**Note:** The HP documentation for your operating system provides details about configuring other system parameters.

3. After you complete the HP-UX installation, install the HP-UX Host Utilities software on the boot LUN.
4. Add the remaining paths to the LUN to the root volume group by entering the following command, which is in the `/opt/Ontap/santools/bin` directory:

```
ontap_config_paths vg00
```

5. Set the host's secondary boot path by entering the following command:

```
setboot -a sec_hw_path_to_lun
```

`sec_hw_path_to_lun` is the secondary hardware path to the boot LUN from the host.

### Example

```
setboot -a 0/6/2/0.1.7.0.0.0.0
```

6. Test the primary boot path by shutting down and rebooting the system:

```
shutdown -r -y 0
```

**Note:** If automatic booting is disabled (set to OFF) for your system, you must boot the system manually.

7. Continue with the installation by proceeding to one of the following tasks:

| If...                         | Follow these instructions...   |
|-------------------------------|--|
| You cannot reboot the system. | Troubleshooting sanboot configurations with PA-RISC systems (LVM only) |
| You can reboot the system     | Testing the secondary boot path for PA-RISC systems (LVM only)         |

## (LVM) Testing the secondary boot path for PA-RISC systems

You should test the secondary boot path before you put a SAN booted configuration into a production environment.

### Steps

1. Shut down the host by entering the following command:

```
shutdown -r -y 0
```

2. At the Main Menu prompt of the BCH, boot the system from the secondary boot path by entering the following command:

```
bo alt
```

3. Stop the secondary boot path and go back to the primary boot path. Shut down the host again:

```
shutdown -r -y 0
```

4. At the Main Menu prompt of the BCH, boot the system from the primary boot path by entering the following command:

```
bo pri
```

## (VxVM) Creating a Veritas-controlled SAN Boot LUN on HP-UX

To use Veritas to configure the HP-UX operating system on the boot LUN, create a SAN boot LUN on HP-UX controlled by Veritas.

### Before you begin

- The Veritas and the ASL should already be installed on the host.
- The LUN that you plan to use for the new SAN boot device should already exist on the storage system and be visible to the host.
- The LUN should have been discovered and be usable by the HP-UX operating system.

### Steps

1. Copy the existing LVM root disk onto the LUN by entering the following command:

```
vxcp_lvmroot
```

The *Veritas Volume Manager 5.0 Administrator's Guide for HP-UX* provides more information about the `vxcp_lvmroot` procedure.

2. Verify that the `vxpfto` and the restore daemon settings are tuned properly.
3. Configure the host's primary and secondary boot path settings using the `setboot` command.

### Related references

[\(Veritas DMP\) Array Support Library and Array Policy Module](#) on page 53

## (VxVM) Creating the boot LUN using a cold install

You can create a SAN boot LUN in a Veritas environment by performing a cold install directly from the HP-UX Core operating system DVD.

### Before you begin

Ensure that the LUN that you plan to use for the new SAN boot device already exists on the storage system and is visible to the host.

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

**Steps**

1. Perform an advanced Ignite-UX installation to the LUN.

**Note:** Ensure that when you select the root disk, Ignite-UX displays only one path to the LUN.

2. After the Ignite-UX installation completes, install the Veritas Storage Foundation on the host.

The host must still see only one path to the LUN.

3. Install the ASL on the host.
4. Install the SAN Toolkit on the host.
5. Make the other paths to the LUN visible to the host.
6. To discover and use the new paths, run the appropriate HP-UX and VxVM commands:

HP-UX commands: `ioscan, insf -e`

VxVM command: `vxdctl enable`

7. Verify that the `vxpfto` and restore daemon settings are tuned properly.
8. Configure the host's primary and secondary boot path settings by entering the following commands:

```
setboot -p primary_path
```

```
setboot -a alternate_path
```

# 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 working.
- 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 `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:

- If you have a cluster, the FC `cfmode` setting is correct and supported by your configuration. The following tools can help you determine which `cfmodes` are supported:
  - N series Interoperability Matrices Web site (accessed and navigated as described in [Websites](#) on page 7).

**Note:** Changing the `cfmode` affects the volume group and zoning on your system. If you must change the `cfmode`, contact technical support.

- FC 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 `fcpc show initiators` command might not display all the host drivers as connected even though they are.
- The `igroups` have the following values for the attributes:
  - The `ostype` attribute must be `hpux`.
  - The `vsa` attribute must be `yes`.
 

**Note:** This attribute only appears in the output displays if it is incorrectly set. When it is correctly set, no references to it appear.
- For LUNs, the following 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

[N series Interoperability Matrices Web site - www.ibm.com/systems/storage/network/interophome.html](http://www.ibm.com/systems/storage/network/interophome.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 rebooting the host

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 on `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.

**Note:** It is best to reboot the host when migrating from a non-ALUA to an ALUA environment.

## Situations affecting volume groups and disk devices

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

These situations include the following:

- If the host reboots while a cluster is in a degraded or takeover state.
- If you change any of the following settings on the storage system:
  - FC cfmode (do not do this unless directed to do so by technical support.)
  - igroup Volume Set Addressing (VSA) setting

## Application I/O failures on Veritas VxVM

Application I/O failures might sometimes occur during storage system cluster 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.
- Verify that the storage system is configured with a supported Data ONTAP cfmode.

## (Veritas DMP) 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 all
controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
2g (2147483648) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
Host      controller          primary   partner   PVlinks
Path      path                /dev/dsk host      controller controller failover
State     type               filename  HBA      port      port      priority
-----
Up        primary           /dev/dsk/c34t0d5 td0  4a                0
Up        primary           /dev/dsk/c35t0d5 td1  5a                not in VG
Up        secondary         /dev/dsk/c30t0d5 td1                5b                not in VG
Up        secondary         /dev/dsk/c31t0d5 td0                4b                not in VG
-----

# ontap_config_paths

Getting information from sanlun...

Adding missing path with:
  vgextend /dev/vg_ntap01 /dev/dsk/c30t0d5
  (output removed)
Moving VG path to the end of the alternate paths with:
  vgreduce /dev/vg_ntap01 /dev/dsk/c34t0d5
  vgextend /dev/vg_ntap01 /dev/dsk/c34t0d5
  (output removed)
Moving VG path to the end of the alternate paths with:
  vgreduce /dev/vg_ntap01 /dev/dsk/c31t0d5
  vgextend /dev/vg_ntap01 /dev/dsk/c31t0d5
  (output removed)
```

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

If your system fails to boot from either the primary or secondary path, you must check the physical connectivity and cabling between the HBA, the switches, the cluster, and the boot path settings.

### Steps

1. If the system does not return to the BCH menu, reset the system to return to the BCH prompt.
2. Boot from the internal system disk by entering the following command:

```
bo path
```

*path* is the hardware path to the internal disk that the system formerly booted from.

The following example uses the hardware path to the internal system disk.

```
bo 0/0/1/1.15.0
```

3. After the system is booted with the internal disk, verify that you correctly translated the primary and secondary device paths to the boot LUN to the correct hardware paths by entering the following commands:

`sanlun lun show -p`: Shows the primary and secondary device paths to the boot LUN.

`ioscan -funC disk`: Shows the hardware paths to the device files.

4. Set the boot path for the boot LUN by entering the following commands:

```
setboot -p primary_path
```

```
setboot -h HA_alternate_path
```

```
setboot -a alternate_path
```

*primary\_path* is the primary path to the boot LUN.

*HA\_alternate\_path* is the secondary path to the boot LUN.

*alternate\_path* is the local disk boot path.

5. Verify the boot path by entering the following command:

```
setboot
```

```
Primary bootpath : 0/3/0/0/0/0.0.1
HA Alternate bootpath : none
Alternate bootpath : 0/3/0/0/0/0.0.0
```

6. Reboot the system to boot from the boot LUN.

## 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

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. Repair the problem that caused the controller to fail, and then reboot the storage system.
2. On the storage console, perform the giveback by entering the following command:  
`cf giveback`
3. 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/clt15d0 /dev/rdisk/clt15d0
disk 14 0/4/0/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c24t0d0 /dev/rdisk/c24t0d0
disk 28 0/4/0/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c26t0d0 /dev/rdisk/c26t0d0
disk 37 0/6/2/0.1.4.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c46t0d0 /dev/rdisk/c46t0d0
```

```
disk 46 0/6/2/0.1.7.0.0.0.0 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c44t0d0 /dev/rdisk/c44t0d0
```

4. 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
controller 1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
2g (2147483648) lun state: GOOD
controller_CF_State: Cluster Enabled Controller Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
host controller primary partner PVlinks
path path /dev/dsk host controller controller failover
state type filename HBA port port priority
-----
up secondary /dev/dsk/c44t0d0 td1 5b 0 (Wrong)
up primary /dev/dsk/c24t0d0 td0 4a 1
up primary /dev/dsk/c46t0d0 td1 5a 2
up secondary /dev/dsk/c26t0d0 td0 4b 3
```

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

```
ontap_config_paths
```

6. 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
```

7. 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 all
controller 1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
2g (2147483648) lun state: GOOD
controller_CF_State: Cluster Enabled Controller Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
```

| host     | controller |                  | primary | partner    | PVlinks    |          |
|----------|------------|------------------|---------|------------|------------|----------|
| path     | path       | /dev/dsk         | host    | controller | controller | path     |
| failover |            |                  |         |            |            |          |
| state    | type       | filename         | HBA     | port       | port       | priority |
| up       | primary    | /dev/dsk/c24t0d0 | td0     | 4a         |            | 0        |
| up       | primary    | /dev/dsk/c46t0d0 | td1     | 5a         |            | 1        |
| up       | secondary  | /dev/dsk/c44t0d0 | td1     |            | 5b         | 2        |
| up       | secondary  | /dev/dsk/c26t0d0 | td0     |            | 4b         | 3        |

## Enable the host to identify more than eight LUNs per target

The `igroup vsa` attribute enables an HP-UX host to identify more than eight LUNs per target HBA. By default, this attribute is set when you create an `igroup` that specifies `hpux` as the operating system type. You must manually set this attribute if you change an existing `igroup` to an HP-UX `igroup` by modifying the `ostype` attribute.

### Steps

1. On the host, make a backup of all the volume groups that contain LUNs.

These backups must contain up-to-date information for both the data within the volume group and the volume group configuration. You should record the LUN volume group membership.

2. Locate all LUNs and volume groups that contain LUNs from a specific pair of storage systems, by entering the following command for each storage system:

```
sanlun lun show -p controller1_name
```

```
sanlun lun show -p controller2_name
```

*controller1\_name*, *controller2\_name* are names of the storage systems.

3. On the host, close the logical volumes.

If a logical volume contains a file system, unmount the file system.

4. On the host, deactivate each volume group by entering the following command:

```
vgchange -a n /dev/vol_group_name
```

*vol\_group\_name* is the path to the volume group.

### Example

The following command deactivates the volume group `/dev/ntap01`:

```
vgchange -a n /dev/ntap01
```

5. On the host, remove the volume group entry from the `/etc/lvmtab` file and associated device files by entering the `vgexport` command:

### Example

```
vgexport /dev/ntap01
```

6. On the storage system, go into diagnostic mode by entering the following command:

```
priv set diag
```

7. On the storage system, set the *vsa* value of the specific *igroup* to *on* or *off* by entering the following command:

```
igroup set igroup vsa value
```

*igroup* is the name of the *igroup* for which you want to change the setting.

*value* is either *on* or *off*. For HP-UX systems, you must always set this value to *on*.

**Note:** Although `igroup set igroup vsa value` command lets you change the *vsa* value of a specific *igroup* and all *igroups* you create in the future, it does not change the value of any other existing *igroups*. You must perform this command on each existing *igroup* to change its attribute value.

8. On the storage system, confirm that the *vsa* value for that *igroup* is now correct by entering the following command:

```
igroup show -v
```

If the attribute is not set correctly for an HP-UX system, the line `Volume Set Addressing: No` appears in the output.

If the attribute is set correctly, no reference to `volume set addressing` or *vsa* appears in the output.

### Example

This example checks the settings for the *igroup* `hostA` and corrects the value for *vsa*.

```
controllerX> igroup show -v hostA
hostA (FCP):
OS Type: hpux
Volume Set Addressing: No
Member: 50:06:0b:00:00:10:a7:00
Member: 50:06:0b:00:00:10:a6:06
controllerX> priv set diag
Warning: These diagnostic commands are for use by Network
Appliance personnel only.
controllerX> igroup set hostA vsa on
Wed Sep 3 21:13:59 GMT [controllerX:
scsitarget.notice:notice]: VSA setting changed on
initiator group "hostA" to "on"
controllerX> igroup show -v hostA
hostA (FCP) (ostype: hpux):
OS Type: hpux
Member: 50:06:0b:00:00:10:a7:00
Member: 50:06:0b:00:00:10:a6:06
controllerX> priv set
```

9. On the host, instruct the host to discover the disks and create device special files for them by entering the following command:

```
ioscan -fnC disk
```

```
ioinit -i
```

**Note:** If the `ioinit -i` command does not create the devices, try the `insf -e` command.

The old disk devices show up as `NO_HW` in the `ioscan` output.

10. On the host, either remove the old device paths manually or reboot the host.

These actions remove the old disk devices and cause the host to discover the new ones. Make sure that the only device special files that are removed are those that disappeared as a result of the `igroup` attribute change.

11. On the host, view the new locations by entering the following commands:

```
vgscan -v
```

```
sanlun lun show -p
```

Correlate the new `sanlun` output with the `sanlun` information from Step 1 to determine which devices go with which volume group.

12. On the host, add the volume group entries back to `/etc/lvmtab` and the associated device files back to the system.

13. For each volume group, perform the following tasks:

- a. Create a new directory by entering the following command:

```
mkdir /dev/vg_name
```

`vg-name` is the volume group name.

- b. Create a group file in the preceding directory by entering the following command:

```
mknod /dev/vg_name/group c 64vg_minor_dev_num
```

- c. Enter the `vgimport` command with a device that represents a primary path (as indicated by the `sanlun lun show -p` command) for the volume group:

```
vgimport /dev/vg_name dev_name_primary_path
```

14. On the host, for each volume group that you exported before changing the `vsa` attribute, activate the newly imported volume group by entering the following command:

```
vgchange -a y /dev/vg_name
```

15. On the host (HP-UX 11iv2), configure multipathing for the newly imported volume groups by running the `ontap_config_paths` utility:

```
ontap_config_paths
```

16. On the host, verify that the paths are now set up correctly by entering the following command:

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

17. Repeat Step 1 to Step 15 on all the hosts that are connected to the storage system.

## 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 `enable_ontap_pvlinks` script

### 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 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/c1t15d0 /dev/rdisk/c1t15d0
disk 1 0/4/0/0.1.5.0.26.0.5 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c31t0d5 /dev/rdisk/c31t0d5
disk 3 0/4/0/0.1.6.0.26.0.5 sdisk NO_HW DEVICE NETAPP LUN
/dev/dsk/c34t0d5 /dev/rdisk/c34t0d5
disk 2 0/6/2/0.1.13.0.26.0.5 sdisk CLAIMED DEVICE NETAPP LUN
/dev/dsk/c30t0d5 /dev/rdisk/c30t0d5
disk 4 0/6/2/0.1.14.0.26.0.5 sdisk NO_HW DEVICE NETAPP LUN
/dev/dsk/c35t0d5 /dev/rdisk/c35t0d5
```

### 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
controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/hp-lv01
2g (2147483648) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
host controller primary partner PVlinks
path path /dev/dsk host controller controller path
failover
state type filename HBA port port priority
-----
up secondary /dev/dsk/c30t0d5 td1 5b 0 (Wrong)
up secondary /dev/dsk/c31t0d5 td0 4b 1 (Wrong)
up primary /dev/dsk/c35t0d5 td1 5a 2
up primary /dev/dsk/c34t0d5 td0 4a not in VG
#
```

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

When adding the `-pin` in the command `sanlun lun show -p`

### 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
controller1:/vol/vol1/hp-lv-lun1 (LUN 1) VG: /dev/vg_ontap01
```

```

2g (2147483648) lun state: GOOD
Controller_CF_State: Cluster Enabled Controller Partner: controller2
Multipath_Policy: A/P
Multipath-provider: Native
-----
host controller primary partner PVlinks
path path /dev/dsk host controller controller path failover
state type filename HBA port port controller path failover
-----
DOWN /dev/dsk/c31t0d5 0
DOWN /dev/dsk/c34t0d5 1
up secondary /dev/dsk/c30t0d5 fcd1 0d 2
up secondary /dev/dsk/c31t0d5 fcd0 0d 3

```

## 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.

The HP-UX 11iv3 February 2007 release does not support ALUA; therefore, the default multipathing policy for all the disk storage is active/active.

**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 `iostscan` 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 Network Appliance, 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 59
    - disabling Native MPIO ALUA for Veritas 19
    - supported with HP-UX 19
  - APM
    - used with Veritas 53
    - available from Symantec 53
    - example of uninstalling 56
    - installing 57
    - obtaining from Symantec 56
    - upgrading 54
    - used with Veritas 53
  - array types
    - displaying with sanlun 66
  - ASL
    - upgrading 54
      - ALUA 59
      - array type 54
      - available from Symantec 53
      - determining the ASL version 55
      - error messages 66
      - example of uninstalling 56
      - installing 57
      - obtaining from Symantec 56
      - used with Veritas 53
- C**
- configurations
    - supported for Host Utilities 17
- E**
- enable\_ontap\_pvlinks
    - installed by Host Utilities 13
  - environments
    - finding information on different Host Utilities environments 15
  - error messages
    - ASL 66
  - example
    - uninstalling ASL, APM 56
- F**
- FCoE
    - converged network adapters 16
    - data center bridging 16
    - Ethernet switch 16
    - traditional FC 16
- H**
- HBA
    - displaying information with sanlun 50
  - Host Utilities
    - contents 13
    - instructions for the different environments 15
    - planning installation 23
    - supported configurations 17
    - using VxVM 61
  - HP-UX 11iv2
    - getting I/O policy information 65
  - HP-UX 11iv3
    - getting I/O policy information 65
  - HP-UX Host Utilities 13
  - HP-UX LVM
    - PVLinks 19
- I**
- I/O policy
    - finding for Veritas 65
  - igroup
    - configuration overview 35
    - uses WWPN 16
  - initiator
    - downloading iSCSI 31
  - installation
    - overview 23
    - planning 23
  - instructions
    - finding information on different environments 15
  - iSCSI protocol
    - downloading initiator 31
    - implementing with Host Utilities 17
    - instructions for that Host Utilities environment 15
- J**
- Jumbo frames
    - not supported with SAN booting 20
- L**
- LUNs

- adding in Veritas environment 61
- configuration overview 35
- creating Veritas boot LUN 69
- displaying with sanlun 47

## LVM

- instructions for that Host Utilities environment 15

## M

man pages

- installed by Host Utilities 13

MPIO

- instructions for that Host Utilities environment 15

multipathing

- using sanlun with Veritas 66

## N

Native MPIO ALUA

- disabling 59

nSANity

- diagnostic utility 13

ntap\_config\_paths

- installed by Host Utilities 13

## P

polling interval

- recommended values 64

protocols

- Fibre Channel (FC) 15

- iSCSI 15

## R

restore policy

- recommended value 64

## S

SAN booting

- advantages 20

- not supported on Jumbo frames 20

san\_version

- installed by Host Utilities 13

sanlun utility

- displaying array types 66

- displaying LUNs 47

- displaying multipathing for Veritas 66

- installed by Host Utilities 13

space reclamation

- command for reclaiming 58

- viewing 58

storage system

- direct-attached 15

- Ethernet network 15

- fabric-attached 15

Symantec

- getting the ASL and APM 56

- provides ASL, APM 53

## T

thin provisioning

- ASL required 58

- viewing 58

## V

Veritas

- restore daemon settings 64

- adding a LUN 61

- APM 53

- ASL 53

- disabling Native MPIO ALUA 59

- displaying multipathing 66

- I/O policy 65

- thin provisioning 58

- using ALUA with HP-UX 19

- using VxVM to display LUN paths 61

- Volume Manager 19

Veritas DMP

- instructions for that Host Utilities environment 15

- SAN booting instructions 20

VxVM

- Veritas Volume Manager 19

- displaying LUN paths 61

## W

WWID 98

WWPN

- getting with sanlun fcp show adapter command 16

- required for igroups 16





NA 210-05376\_A0, Printed in USA

GA32-0764-02

