



**IBM System Storage N series
AIX Host Utilities 5.1 Installation and Setup Guide**

Contents

About this guide	7
Supported features	7
Getting information, help, and services	7
Before you call	8
Using the documentation	8
Web sites	8
Accessing online technical support	9
Hardware service and support	9
Supported servers and operating systems	9
Firmware updates	9
How to send your comments	9
Changes to this document: January 2011	11
The AIX Host Utilities	13
Supported Host Utilities environments and protocols	14
How to find instructions for your AIX Host Utilities environment	18
Quick start checklist for planning the Host Utilities installation	19
(FC) Setting the media type to loop for direct-attached	
configurations	23
The AIX Host Utilities installation	25
Getting the software	26
Using SMIT to install the Host Utilities software	28
Using a command line to install the Host Utilities software	33
Values for the queue_depth and num_cmd_elems parameters	35
(FC) AIX fc_err_recov parameter settings	36
Information on removing the AIX Host Utilities	37
Before you uninstall previous versions of software	37
Uninstalling the Host Utilities software using SMIT	37
Using a command line to uninstall the Host Utilities software	41
(iSCSI) Additional configuration for iSCSI environments	43
(iSCSI) Changing the software initiator's iSCSI node name	43
(iSCSI) Prepare the HBA for configuration	44
(iSCSI) Configuring the HBA	45

(iSCSI) Adding information to the targets file	47
(iSCSI) Adding targets	48
(iSCSI) CHAP authentication	49
(iSCSI) Data ONTAP upgrades may affect CHAP configuration	49
(Veritas DMP) Configuration requirements for Veritas Storage	
Foundation environments	51
(Veritas DMP) The Array Support Library and the Array Policy Module	51
(Veritas DMP) Information provided by the ASL	52
(Veritas DMP) Information on installing and upgrading the ASL and APM	53
(Veritas DMP) Removing the ASL and APM	57
(Veritas DMP) Using sanlun to obtain multipathing information	58
(Veritas DMP) Using VxVM to display available paths	58
(Veritas DMP) The Veritas DMP restore daemon requirements	59
(Veritas DMP) Setting the restore daemon interval	59
(Veritas DMP) Information on ASL error messages	60
Overview of LUN configuration and management	63
(AIX MPIO, PowerVM) Discovering and Managing LUNs	64
(AIX MPIO without PowerVM) Discovering and configuring LUNs with AIX	65
(PowerVM VSCSI) Discovering and configuring LUNs	66
(PowerVM NPIV) Discovering and configuring LUNs	70
(Veritas DMP) Discovering and configuring LUNs with Veritas	72
(Veritas DMP) Discovering LUNs on a host using VxVM and DMP	72
(Veritas DMP) Methods for adding a disk to a disk group	73
(Veritas DMP) Methods for creating a volume	74
(Veritas DMP) Methods for accessing storage on a Veritas volume	75
The sanlun utility	77
Displaying host LUN information with sanlun	77
Explanation of the sanlun lun show -p output	78
Displaying host HBA information with sanlun	79
SAN boot LUNs in an AIX MPIO environment	81
(AIX MPIO) Important notes about creating a SAN boot LUN	81
(AIX MPIO, FC) Creating a SAN boot LUN in an FC environment	82
(AIX MPIO, FC) Cloning the root volume group of the AIX operating system	86
(AIX MPIO, iSCSI) Creating a SAN boot LUN in an iSCSI environment	88

(AIX MPIO) Updating a SAN boot LUN	99
Features supported by the Host Utilities	103
(AIX MPIO) Multiple protocols supported on a single host	104
HBAs and the AIX Host Utilities	104
Multipathing and the AIX Host Utilities	104
Dynamic LPAR is available	104
GPFS support	105
Live Partition Mobility from PowerVM 1.5.2	105
(FC) ALUA support with certain versions of Data ONTAP	105
(FC) Specific AIX Host Utilities environments that support ALUA	106
SAN booting and the Host Utilities	107
Support for non-English versions of AIX operating systems	107
Troubleshooting	109
Check the Release Notes	109
If you cannot resolve your problem	109
About the troubleshooting sections that follow	110
Possible connectivity problems	110
Storage system checklist	111
Host checklist	111
Useful commands and utilities for troubleshooting	112
Check the version of your host operating system	113
Check the version of the prerequisite filesets	113
Confirm that the HBA is supported	114
Check the LUN characteristics	114
Check the MPIO disk path information	115
Reinstalling if the installation fails due to system requirements	116
FC troubleshooting tips	116
(FC) Check the ALUA settings on the storage system	116
iSCSI troubleshooting tips	117
(MPIO, iSCSI) Verify information about the HBA	117
(MPIO, iSCSI) Verify that the HBA is logged into the storage system	118
(MPIO, iSCSI) Verify the adapter's configuration	118
Basic operations take a long time to complete	119
Installing the nSANity data collection program	120
Where to find more information	123
Copyright information	125

Trademark information	127
Index	129

About this guide

Here you can learn what this document describes and who it is intended for, what special terminology is used in the document, what command, keyboard, and typographic conventions this document uses to convey information, and other details about finding and using information.

This guide describes how to configure an AIX host in a storage area network (SAN) environment. The host accesses data on a storage system that runs Data ONTAP software with the Fibre Channel (FC) or iSCSI protocol.

See the IBM support page more information about using these hosts with IBM N series storage systems.

Next topics

Supported features on page 7

Getting information, help, and services on page 7

How to send your comments on page 9

Related information

IBM Support page - www.ibm.com/storage/support/nas/

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 at the following Web site:

www.ibm.com/storage/support/nas/

A listing of currently available N series products and features can be found at the following Web site:

www.ibm.com/storage/nas/

Getting information, help, and services

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.

Next topics

- [*Before you call* on page 8](#)
- [*Using the documentation* on page 8](#)
- [*Web sites* on page 8](#)
- [*Accessing online technical support* on page 9](#)
- [*Hardware service and support* on page 9](#)
- [*Supported servers and operating systems* on page 9](#)
- [*Firmware updates* on page 9](#)

Before you call

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

- Check all cables to make sure that they are connected properly.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

Using the documentation

Information about N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Data ONTAP software publications are available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Web sites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- For NAS product information, go to the following Web site:
www.ibm.com/storage/nas/
- For NAS support information, go to the following Web site:
www.ibm.com/storage/support/nas/
- For AutoSupport information, go to the following Web site:
www.ibm.com/storage/support/nas/
- For the latest version of publications, go to the following Web site:
www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

www.ibm.com/storage/support/nas/

Hardware service and support

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

www.ibm.com/planetwide/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, follow the link to the Interoperability Matrices from the following Web site:

www.ibm.com/systems/storage/network/interophome.html

Firmware updates

As with all devices, it is recommended that you run the latest level of firmware, which can be downloaded by visiting the following Web site:

www.ibm.com/storage/support/nas/

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support. See the *Data ONTAP Upgrade Guide* for your version of Data ONTAP for more information on updating firmware.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com. Be sure to include the following:

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

Changes to this document: January 2011

Several changes have been made to this document since it was published for the AIX Host Utilities 5.0 release.

This document has been updated based on the following major changes:

- Information on the diagnostic utilities has been removed.
Previously, the Host Utilities included the `controller_info`, `filer_info`, `aix_info`, and `switch_info` utilities. Starting with the AIX 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, you can download the latest version of nSANity. You should only download and run this program if Technical Support asks you to.
- Information on using Veritas Storage Foundation 5.1 with the Host Utilities has been added.
This version of Veritas Storage Foundation includes:
 - The Array Support Library (ASL) and the Array Policy Module (APM), so you do not need to install them separately.
 - Support for ALUA.
 - Support for creating a SAN boot LUN.

Note: For instructions on creating a SAN boot LUN in a Veritas environment, see the Symantec documentation for Veritas Storage Foundation 5.1. The Veritas Storage Foundation DocCentral is available at <http://sfdoccentral.symantec.com/>.

This version of the Host Utilities supports Veritas Storage Foundation 5.x. It does not support Veritas Storage Foundation 4.x.

Note: To determine which versions of Veritas Storage Foundation are supported with the AIX Host Utilities 5.1, see the IBM NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.

The AIX Host Utilities

The Host Utilities provide software programs and documentation that you can use to connect your AIX host to storage systems running Data ONTAP. The software is available as a single compressed file.

The Host Utilities include the following components:

- The SAN Toolkit

Note: This toolkit is common across all configurations and protocols of the 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.

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

- **(AIX MPIO, FC)** `dotpaths` utility, which lets you set path priorities in MPIO environments running the FC protocol that are not using ALUA.
- The `san_version` command, which displays the versions of the Host Utilities.
- `sanlun` utility. This utility displays information about LUNs on the storage system that are available to this host. In PowerVM environments, you must be running N_Port ID Virtualization (NPIV) in an FC environment if you want to run the `sanlun` utility on an VIO client. If you are running PowerVM with Virtual SCSI (VSCSI), you cannot run Host Utilities tools such as the `sanlun` utility on VIO clients. In VSCSI deployments, the tools only work on the VIO server.

For details on using these tools, see the man pages.

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. You should download, install, and execute it only when requested to do so by technical support.

- Documentation

The documentation provides information on 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.

- *Host Settings Affected by AIX Host Utilities*
- *Quick Command Reference*

You can download the documentation when you download the Host Utilities software.

Next topics

[Supported Host Utilities environments and protocols](#) on page 14

[How to find instructions for your AIX Host Utilities environment](#) on page 18

Supported Host Utilities environments and protocols

The Host Utilities support several AIX environments. These environments include MPIO, PowerVM, and Veritas Storage Foundation.

For details on which environments the Host Utilities support, check the online *[NAS Interoperability Matrices Web site](#)*.

The following table summarizes key aspects of the primary Host Utilities environments for hosts running AIX.

AIX Environment	Notes
AIX MPIO (Native OS)	<p>This environment works with features provided by the AIX operating system to provide volume management and multipathing.</p> <ul style="list-style-type: none"> • Multipathing: AIX's Multiple Path I/O (MPIO) with IBM's default Path Control Module (PCM). • Volume management: AIX Logical Volume Manager (LVM). • Protocols: FC and iSCSI. • ALUA (Asymmetrical Logical Unit Access): You can use either ALUA or the <code>dotpaths</code> utility provided with the Host Utilities to set the path priorities. • Installation software: Install the software packages that are placed in the following two directories when you uncompress the Host Utilities compressed file: <ul style="list-style-type: none"> • MPIO directory: <code>Ontap.MPIO_Host_Utils_Kit</code> • SAN_Tool_Kit directory: <code>Ontap.SAN_toolkit</code> <p>Note: The SAN Toolkit software package is common to all AIX environments and must be installed in addition to the software package for your AIX environment.</p> <ul style="list-style-type: none"> • Setup issues: None. • LUN management: <ul style="list-style-type: none"> • To discover LUNs, use the <code>cfgmgr</code> command. • To manage the LUNs and volumes, use LVM. • To set LUN path priorities: <ul style="list-style-type: none"> • For FC environments that support ALUA and have it enabled, ALUA automatically sets the paths. • For FC environments that are not using ALUA, use the <code>dotpaths</code> utility provided by the Host Utilities. • SAN booting is supported

AIX Environment	Notes
PowerVM	<p>PowerVM uses the Virtual I/O (VIO) Server/Client. This environment runs on top of the MPIO environment and, except where noted, uses the same AIX features.</p> <ul style="list-style-type: none"> Special notes about PowerVM environments: <ul style="list-style-type: none"> You cannot use Veritas Dynamic Multipathing (DMP) on a VIO Server. You can run PowerVM with VSCSI and NPIV. You must run as <code>padmin</code> and use VIO commands. Some Host Utilities require you to run as <code>root</code>. To enable you to become root while logged in as <code>padmin</code> in a PowerVM VSCSI environment, use the <code>oem_setup_env</code> command. Multipathing: MPIO with IBM's default PCM. Volume management: AIX LVM. Protocols: FC and iSCSI. ALUA: You can use either ALUA or the <code>dotpaths</code> utility provided with the Host Utilities to set the path priorities. Installation software: Install the software packages that are placed in the following two directories when you uncompress the Host Utilities compressed file: <ul style="list-style-type: none"> MPIO directory: <code>Ontap.MPIO_Host_Utilities_Kit</code> SAN_Tool_Kit directory: <code>Ontap.SAN_toolkit</code> <p>Note: The SAN Toolkit software package is common to all AIX environments and must be installed in addition to the software package for your AIX environment.</p> Setup issues: None. LUN management: <ul style="list-style-type: none"> To discover LUNs, use the <code>cfgmgr</code> command. To manage the LUNs and volumes, use LVM. To set LUN path priorities: <ul style="list-style-type: none"> For FC environments that support ALUA and have it enabled, ALUA automatically sets the paths. For FC environments that are not using ALUA, use the <code>dotpaths</code> utility provided by the Host Utilities. If you have virtual SCSI LUNs set up, the <code>sanlun</code> utility does not recognize them. This is because the LUNs are identified as VID "IBM" and PID "VDASD". SAN booting is supported.

AIX Environment	Notes
Veritas DMP	<p>This environment uses Veritas Storage Foundation to provide multipathing and volume management.</p> <p>Note: Check the NAS Interoperability Matrices Web site to identify the supported versions of Veritas Storage Foundation.</p> <ul style="list-style-type: none"> • You cannot use the Veritas DMP on a VIO Server. • Multipathing: Veritas DMP • Volume management: Veritas Volume Manager (VxVM) • Protocols: FC • ALUA: Veritas Storage Foundation 5.1 supports ALUA. • Installation software: Install both the host settings software package and the SAN Toolkit software package that are placed in the following two directories when you uncompress the Host Utilities file: <ul style="list-style-type: none"> • SAN directory: <code>Ontap.FCP_Host_Utils_Kit</code> • SAN_Tool_Kit directory: <code>Ontap.SAN_toolkit</code> <p>Note: The SAN Toolkit software package is common to all AIX environments and must be installed in addition to the software package for your AIX environment.</p> <ul style="list-style-type: none"> • Setup issues: <ul style="list-style-type: none"> • The Symantec Array Support Library (ASL) and Array Policy Module (APM) must be installed. If you are using Veritas Storage Foundation 5.1, these modules are automatically installed when you install Veritas Storage Foundation. • LUN management: <ul style="list-style-type: none"> • To discover LUNs, use the <code>cfgmgr</code> command followed by the <code>vxldctl enable</code> command . • To manage the LUNs and volumes, use VxVM. • SAN booting is supported with Veritas Storage Foundation 5.1. See the Symantec documentation (Veritas Storage Foundation DocCentral at http://sfdoccentral.symantec.com/) for information on setting up SAN booting. You do not need any special instructions for setting up a SAN boot LUN with a storage system running Data ONTAP.

How to find instructions for your AIX Host Utilities environment

Many instructions in this manual apply to all the environments supported by the Host Utilities. In some cases, though, commands or configuration information varies based on your environment. To make finding information easier, this guide places a qualifier, such as "**Veritas DMP**," in the title if a section applies only to a specific Host Utilities environment.

That way you can quickly determine whether a section applies to your Host Utilities environment and skip the sections that do not apply.

If the information applies to all supported AIX Host Utilities environments, there is no qualifier in the title.

This guide uses the following qualifiers to identify the different AIX Host Utilities environments:

Qualifier	The section that follows applies to
(Veritas DMP)	Environments using Veritas DMP as the multipathing solution.
(PowerVM)	Environments using PowerVM VIO servers. Unless otherwise specified, the information for MPIO environments also applies to PowerVM environments.
(PowerVM NPIV)	Environments using PowerVM VIO servers with NPIV.
(PowerVM VSCSI)	Environments using PowerVM VIO servers with Virtual SCSI (VSCSI).
(AIX MPIO)	Environments using AIX MPIO as the multipathing solution. Currently, all AIX MPIO environments use native drivers. Unless otherwise specified, the information for "AIX MPIO" environments also applies to Host Utilities environments that use PowerVM.
(AIX MPIO, FC)	AIX MPIO environments using the FC protocol. Unless otherwise specified, the information for "AIX MPIO, FC" environments also applies to PowerVM environments.
(MPIO, iSCSI)	AIX MPIO environments using the iSCSI protocol. Unless otherwise specified, the information for "AIX MPIO, iSCSI" environments also applies to PowerVM environments.
(FC)	Environments using the Fibre Channel protocol.
(iSCSI)	Environments using the iSCSI protocol.

There is also information about using the Host Utilities in an AIX environment in the *Release Notes* and the AIX Host Utilities reference documentation. You can download all the Host Utilities documentation.

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 following checklist provides a high-level overview of these tasks.

If you are an experienced AIX 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 may 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 [NAS Interoperability Matrices Web site](#) for the most current information about system requirements. This includes:
 - Host operating system version, technology levels, and appropriate updates.
 - **(PowerVM)** If you are using PowerVM, make sure you have that environment set up.
 - HBAs and drivers, or software initiators, model and version.

Note: (AIX MPIO, FC) If you are using the FC protocol and a direct-attached configuration, the HBA works with the default driver settings, but you must configure the target HBA on the storage system for a media type of loop.

```
4a: ONLINE <ADAPTER UP> PTP Fabric
      host address
      portname 50:0a:09:81:97:c9:7e:5c nodename 50:0a:
09:80:87:c9:7e:5c
      mediatype loop speed auto
```

- **(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 (Veritas Storage Foundation 5.1 includes these modules in its installation). You must also set the Veritas DMP restore daemon interval to 60.

Note: If you have Veritas Storage Foundation 5.0, you must download and manually install the ASL and APM. Make sure you download the correct versions of the ASL and APM for your Veritas environment. You can get the ASL and APM from the Symantec Web site.

- Volume management and multipathing.

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.
- **(FC active/active environments)** Set up to work with ALUA, if it is supported with your version of Data ONTAP and you are using the FC protocol. With Data ONTAP 8.0 or later, ALUA is enabled by default when you create an igrup in an FC environment. Earlier versions of Data ONTAP require that you manually enable ALUA. If you have an MPIO environment and you did not enable ALUA, you can use the `dotpaths` utility that comes with the Host Utilities.

In Veritas environments, the ASL detects whether ALUA is enabled. It sets the array type to NetApp-A/A for non-ALUA environments and to ALUA for ALUA environments. Veritas Storage Foundation 5.1 supports ALUA.

Note: ALUA is not supported with iSCSI.

- Set up with working volumes and qtrees (if desired).

3. **(FC)** 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 Service and Support Web site .

Uncompress the file and get the SAN Toolkit software package (`Ontap.SAN_toolkit`) as well as the host settings file for your environment (AIX MPIO, PowerVM: MPIO/ `Ontap.MPIO_Host_Utils_Kit`; Veritas DMP: SAN/ `Ontap.FCP_Host_Utils_Kit`; or single-path iSCSI: SAN/ `Ontap.iSCSI_Host_Utils_Kit`).

2. If you are upgrading the Host Utilities, remove any currently installed version of the Host Utilities. You can uninstall the Host Utilities using either the AIX System Management Interface Tool (SMIT) or the command `installp -uFileSetName` . If you are using PowerVM with

VSCSI, you must first log in as `padmin` and enter the `oem_setup_env` command to become root before you can uninstall the Host Utilities.

3. From the directory that contains the extracted software packages, use either the SMIT or the command `installp -u FileSetName` to install the host settings software for your environment and the SAN Toolkit software. You must be logged in as root to do this.
4. Verify that the HBA initiator queue depth is appropriate for your system. You can use the `lsattr -El hdiskx` command to check the queue depth.

(iSCSI) Task 3: Configure the iSCSI protocol

1. Record the host's iSCSI node name.

2. Set up the initiator.

If you have a software initiator:

- Generate a name for the software initiator's iSCSI node. The default name that AIX generates is for demo purposes only.

If you have an HBA:

- Prepare the HBA by either removing the protocol device if it is configured or changing the database.
- Configure the HBA.

3. Add information to the targets file.

4. Set up target discovery.

5. **(Optional)** Set up Challenge Handshake Authentication Protocol (CHAP) on the host and the storage system.

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. If you have Data ONTAP 8.0 or later, ALUA is automatically enabled when you create an igroup in an environment using the FC protocol.

You can use the command `igroup show -v igroup_name` to see if ALUA is enabled. If it is not enabled, you can use the command `igroup set igroup_name alua yes` to enable it.

2. **(FC)** Set the path priorities. If you have an MPIO environment running the FC protocol and you are not using ALUA, you can use the `dotpaths` utility that comes with the Host Utilities to set the path priorities.

3. Discover the new LUNs.

- **(AIX MPIO, Veritas)** Use the `cfgmgr` command. For Veritas you must also use the `vxdctl enable` command.
- **(PowerVM)** Use the `cfgdev` command.

4. **(PowerVM)** If you have LUNs presented to a VIO server from multiple third-party storage vendors, make sure that all the LUNs use the same maximum transfer size. Use the `lsattr -El <disk> -a max_transfer` command.
5. Display information about the LUNs. You can use the `sanlun` command to display information about the LUNs and the HBAs.

(FC) Setting the media type to loop for direct-attached configurations

If you use a direct-attached, FC configuration, the HBA in the host can use the default settings, but you must set the target HBA media type to loop. You can set the media type by running the `fcp` command on the storage system.

Before you begin

Make sure that you do not have clustering licensed or the Interconnect Card installed. Remove both if they exist on your storage system.

Note: If you have enough initiators in the host, you can set up a direct-attached configuration that is connected to a clustered controller.

Steps

1. Stop each adapter used in the direct-attached configuration by entering the following command:

```
fcp config adapter down
```

adapter is the name of the FC target adapter.

2. Set the media type to loop for each adapter by entering the following command:

```
fcp config adapter mediatype loop
```

3. Restart each adapter by entering the following command:

```
fcp config adapter up
```

4. Reboot the storage system and the host.

5. Verify the adapter ports are online by entering the following command:

```
fcp show adapter
```

6. Verify that the media type is now correctly set on each adapter port by entering the following command:

```
fcp config
```


The AIX Host Utilities installation

The AIX Host Utilities software is provided in a single, compressed file that contains the software packages for AIX MPIO and PowerVM environments, Veritas environments, single-path iSCSI environments, and the SAN Toolkit. You must install both the host settings software package for your environment and the SAN Toolkit package.

You can perform the installations using either SMIT or the `installp` command.

Installing the Host Utilities requires that you install two software packages:

- The SAN Toolkit software package.

The SAN Toolkit software package contains the `sanlun` utility and other tools. You install the same software package for all environments of the Host Utilities. All environments use the same SAN Toolkit software package.

- The host settings software package for your environment.

The host settings are the Object Data Manager (ODM) stanza LPP packages that contain the device definitions for your multipathing environment. There is a separate software package for MPIO, Veritas DMP, and single-path iSCSI environments.

Note: Direct-attached configurations and switch configurations that have multiple paths should use the MPIO software package.

When you uncompress the AIX Host Utilities file, it creates a directory within its current directory that has three subdirectories. The subdirectories contain the Host Utilities software packages for the MPIO host settings, the Veritas DMP host settings, and the SAN Toolkit.

Note: It is a good practice to uninstall any currently installed version of the Host Utilities before you install a new version or even switch Host Utilities environments, such as going from MPIO to Veritas. By performing a clean upgrade, you can avoid possible problems in the future.

Your cluster failover mode (cfmode) must be set to single-image mode. Any time you change the cfmode, you must reconfigure your host to discover the new paths to the disk.

Next topics

[*Getting the software*](#) on page 26

[*Using SMIT to install the Host Utilities software*](#) on page 28

[*Using a command line to install the Host Utilities software*](#) on page 33

[*Values for the queue_depth and num_cmd_elems parameters*](#) on page 35

[*\(FC\) AIX fc_err_recov parameter settings*](#) on page 36

Getting the software

You can download the compressed file containing the Host Utilities software packages or you can install it from the physical media. Once you have the file, you must uncompress it to get the two software packages you need to install the Host Utilities.

About this task

You can install the product software either from the physical media kit or from software updates available for download. Downloads are available only to entitled IBM N series customers who have completed the registration process on the IBM NAS support web site at www.ibm.com/storage/support/nas/.

Steps

1. Check the publication matrix page at www.ibm.com/systems/storage/network/interophome.html for important alerts, news, interoperability details, and other information about the product before beginning the installation.
 - If you are installing the software from the Physical media kit, insert the CD-ROM into your host machine. You are now ready to proceed to the instructions for installing the software.
 - If you are installing the software from the Software updates available for download, go to the IBM NAS support web site at www.ibm.com/storage/support/nas/. Now follow the steps (2-4) listed below.

2. Sign in with your IBM ID and password.

If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration:

- a. Click the Register link.
- b. Follow the online instructions.
- c. Sign in.

3. Select the N series software you want to download, and then select the Download view.

4. Use the “Software Packages ...” link on the Web page presented and follow the online instructions to download the software.

Note: If you download the file to a machine other than the AIX host, you must move the file or a copy of it to the AIX host.

5. On the AIX host, make sure you are logged in as root.

Note: (PowerVM) In PowerVM environments with VSCSI, log in as `padmin` and then enter the `oem_setup_env` command to become root.

6. Go to the directory where you placed the Host Utilities file and uncompress it.

You can use the `zcat` and `tar` commands to uncompress the file and extract the software:

```
zcat ibm_aix_host_utilities_5.1.tar.Z | tar -xvf -
```

Note: The actual name of the file containing the Host Utilities software may differ from the file name shown in these examples.

When you uncompress this file, it creates the directory `ibm_aix_host_utilities_5.1` with the following three subdirectories:

- **SAN_Tool_Kit** (Required for all environments)

This directory contains the SAN Toolkit software package (`Ontap.SAN_toolkit`), which applies to all AIX Host Utilities environments and works with all multipathing solutions.

- **MPIO**

This directory contains the host settings software package

(`Ontap.MPIO_Host_Utils_Kit`) for supported Host Utilities environments using AIX MPIO for multipathing.

Both AIX MPIO and PowerVM environments must install this software package.

- **SAN**

This directory contains two host settings software packages:

- If you are using Veritas DMP as your multipathing solution, you must install the `Ontap.FCP_Host_Utils_Kit` software package.
- If you are using single-path iSCSI, you must install the `Ontap.iSCSI_Host_Utils_Kit` software package.

To install the Host Utilities, you must install both the host settings software package for your multipathing solution and the SAN Toolkit software.

Example

The following example shows the type of output that executing the `zcat` and `tar` commands produces.

```
#-> zcat ibm_aix_host_utilities_5.1.tar.Z | tar xvf -
x ibm_aix_host_utilities_5.1
x ibm_aix_host_utilities_5.1/SAN_Tool_Kit
x ibm_aix_host_utilities_5.1/SAN_Tool_Kit/Ontap.SAN_toolkit, 1343488 bytes,
2624 tape blocks
x ibm_aix_host_utilities_5.1/SAN
x ibm_aix_host_utilities_5.1/SAN/Ontap.fcp_Host_Utils_Kit, 24576 bytes, 48
tape
blocks
x ibm_aix_host_utilities_5.1/SAN/Ontap.iscsi_Host_Utils_Kit, 20480 bytes, 40
tape
blocks
x ibm_aix_host_utilities_5.1/MPIO
x ibm_aix_host_utilities_5.1/MPIO/Ontap.mpio_Host_Utils_Kit, 184320 bytes,
360 tape
block
```

Using SMIT to install the Host Utilities software

You can use SMIT to install the Host Utilities software for your environment.

Before you begin

Make sure you have the host settings software package for your environment and the SAN Toolkit package that were contained in the compressed file you downloaded or got from the physical media.

In addition, it is a good practice to check the Host Utilities *Release Notes* to see if there have been any changes or new recommendations for installing and using the Host Utilities since this installation guide was produced.

About this task

This installation example installs the host settings software package first. You must repeat the steps to install the SAN Toolkit software package.

(PowerVM) If you have a PowerVM environment, you must install the SAN Toolkit on each client. If you are running NPIV in an FC environment, you can then run the `sanlun` utility on the VIO client. If you have a PowerVM VSCSI, the `sanlun` utility only runs on the VIO server; not the VIO clients.

Steps

1. Log in to the host as root.

(PowerVM) In PowerVM VSCSI environments, you must log into the host as `padmin` and then enter the `oem_setup_env` command to become root.

Example

(PowerVM) Entering the following command from a VIO server that uses VSCSI enables you to become root.

```
vioserver$ oem_setup_env
```

2. Start SMIT by entering the following command:

```
smitty install
```

3. At the screen that appears, select the option titled Install and Update Software.

Example

When you start SMIT, it displays the following screen. On this screen, the Install and Update Software option is the first menu option.

```
# smitty install
+-----+
--+
Software Installation and Maintenance
```

```

Move cursor to desired item and press Enter.
  Install and Update Software
  List Software and Related Information
  Software Maintenance and Utilities
  Software Service Management
  Network Installation Management
  EZ NIM (Easy NIM Tool)
  System Backup Manager
  Alternate Disk Installation
  EFIX Management

```

```

F1=Help      F2=Refresh      F3=Cancel      Esc+8=Image
Esc+9=Shell  Esc+0=Exit      Enter=Do
+-----+
---+

```

4. At the next screen, select the Install Software menu option.

Example

The following is an example of the Install and Update Software screen.

```

+-----+
---+
Install and Update Software
Move cursor to desired item and press Enter.
  Install Software
  Update Installed Software to Latest Level (Update All)
  Install Software Bundle
  Update Software by Fix (APAR)
  Install and Update from ALL Available Software

F1=Help      F2=Refresh      F3=Cancel      Esc+8=Image
Esc+9=Shell  Esc+0=Exit      Enter=Do
+-----+
---+

```

5. At the Install Software screen, specify the location of the software.

For example, when you're installing from the host machine, you would enter the path to the software package. If you use MPIO for multipathing, you might enter something similar to:

```
/var/tmp/ibm_aix_host_utilities_5.1/MPIO/Ontap.MPIO_Host_UtilsKit
```

If you're installing from the CD-ROM, enter the CD-ROM drive.

Example

(AIX MPIO) The following is an example of entering the path to the MPIO software package when you are at the Install Software screen.

```

+-----+
Install Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

```

[Entry Fields]

```

* INPUT device / directory for software
[/var/tmp/ibm_aix_host_utilities_5.1/MPIO/Ontap.MPIO_Host_UtilsKit]
+
F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  Esc+6=Command  Esc+7>Edit
Esc+8=Image
Esc+9=Shell  Esc+0=Exit     Enter=Do
+-----+

```

After you enter path to the software package, SMIT displays the following screen. This is the screen where you enter name of the software package and press Enter.

```

+-----+
Install Software

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

[Entry Fields]
* INPUT device / directory for software
[/var/tmp/ibm_aix_host_utilities_5.1/MPIO/Ontap.MPIO_Host_UtilsKit]
+
* SOFTWARE to install
[Ontap.MPIO_Host_UtilsKit] +

PREVIEW only? (install operation will NOT occur)      no      +
COMMIT software updates?                            yes
+
SAVE replaced files?                                no      +
AUTOMATICALLY install requisite software?          yes      +
EXTEND file systems if space needed?                yes      +
OVERWRITE same or newer versions?                  no      +
VERIFY install and check file sizes?              no      +
Include corresponding LANGUAGE filesets?          yes      +
DETAILED output?                                    no      +
Process multiple volumes?                          yes      +
ACCEPT new license agreements?                  no      +
Preview new LICENSE agreements?                  no      +
+-----+
F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  Esc+6=Command  Esc+7>Edit
Esc+8=Image
Esc+9=Shell  Esc+0=Exit     Enter=Do
+-----+

```

6. SMIT displays a warning, asking you to confirm that you want to install the software. Press Enter to continue.

Example

The following is an example of the screen that SMIT displays.

```

*****x

```

```

x                               ARE YOU SURE?
x
x Continuing may delete information you may want
x to keep. This is your last chance to stop
x before continuing.
x     Press Enter to continue.
x     Press Cancel to return to the application.
x
x F1=Help          F2=Refresh
F3=Cancel
x Esx Esc+8=Image Esc+0=Exit   Enter=Do
x*****

```

7. SMIT displays information about the installation.

Note: By default, the host settings software installs the filesets for both FC and iSCSI. This is not a problem. If you are using only one protocol, ignore the other protocol.

Example

The following is the type of output you see when an installation using the MPIO fileset successfully completes.

```

COMMAND STATUS

Command: running      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

Ontap.MPIO_Host_Utils_Kit.pcmodm 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.iscsi 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.fcp 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.config 5.1.0.0
. . . . . ** Copyright notice for Ontap.MPIO_Host_Utils_Kit
**. . . . .
(C) Copyright NetApp, Inc. 2003-2011.
All rights reserved.
. . . . . ** End of copyright notice for Ontap.MPIO_Host_Utils_Kit
**. . . . .

Finished processing all filesets. (Total time: 5 secs).

0503-409 installlp: bosboot verification starting...
installlp: bosboot verification completed.
0503-408 installlp: bosboot process starting...

bosboot: Boot image is 23825 512 byte blocks.
0503-292 This update will not fully take effect until after a
          system reboot.

* * * A T T E N T I O N * * *
System boot image has been updated. You should reboot the
system as soon as possible to properly integrate the changes
and to avoid disruption of current functionality.

installlp: bosboot process completed.
+-----+

```

```
Summaries:
+-----+
Installation Summary
+-----+
Name           Level      Part      Event
Result

+---+
Ontap.MPIO_Host_Utils_Kit.pcmod 5.1.0.0 USR APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.iscsi 5.1.0.0 USR APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.fcp 5.1.0.0 USR APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.conf 5.1.0.0 USR APPLY SUCCESS

COMMAND STATUS

Command: running      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

Ontap.MPIO_Host_Utils_Kit.config 5.1.0.0

. . . . . ** Copyright notice for Ontap.MPIO_Host_Utils_Kit
** . . . . .
(C) Copyright NetApp, Inc. 2003-2011.
All rights reserved.
. . . . . ** End of copyright notice for Ontap.MPIO_Host_Utils_Kit
**. . . .

Finished processing all filesets. (Total time: 5 secs).

0503-409 installlp: bosboot verification starting...
installlp: bosboot verification completed.
0503-408 installlp: bosboot process starting...

bosboot: Boot image is 23825 512 byte blocks.
0503-292 This update will not fully take effect until after a system
reboot.

installlp: bosboot process completed.
+-----+
Summaries:
+-----+
Installation Summary
+-----+



* * * A T T E N T I O N * * *
System boot image has been updated. You should reboot the
system as soon as possible to properly integrate the changes
and to avoid disruption of current functionality.

+-----+
Summaries:
+-----+
Installation Summary
```

Name	Level	Part	Event
Result			
Ontap.MPIO_Host_Utils_Kit.pcmod	5.1.0.0	USR	APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.iscsi	5.1.0.0	USR	APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.fcp	5.1.0.0	USR	APPLY SUCCESS
Ontap.MPIO_Host_Utils_Kit.conf1	5.1.0.0	USR	APPLY SUCCESS

F1=Help F2=Refresh F3=Cancel Esc
 +6=Command Esx Esc+8=Image Esc+9=Shell Esc+0=Exit

- Now that you have installed the host settings software package, you should install the SAN Toolkit software package: Ontap.SAN_toolkit. Repeat Steps 2-8 to install the toolkit.

Note: In a PowerVM environment, you should also install the SAN Toolkit on each client.

You can check the installation by looking in the SMIT log file (smit.log and smit.script). These files are in the SMIT log directory (\$HOME).

After you finish

To complete the installation, you must configure any parameters required for your environment. For example, if you are using iSCSI, you must configure the initiator on the host. You also must set up LUNs and have the host discover them.

Using a command line to install the Host Utilities software

You can use command line to install the Host Utilities software for your environment.

Before you begin

Make sure you have the host settings software package for your environment and the SAN Toolkit package from the Host Utilities compressed file you downloaded or got from the physical media.

In addition, it is a good practice to check the Host Utilities *Release Notes* to see if there have been any changes or new recommendations for installing and using the Host Utilities since this installation guide was produced.

About this task

(PowerVM) If you have a PowerVM environment, you must install the SAN Toolkit on each client. If you are running NPIV in an FC environment, you can then run the `sanlun` utility on the VIO client. Without NPIV, you can only run the `sanlun` utility on the server; not the VIO clients.

Steps

- Log in to the host as root.

(PowerVM) In PowerVM environments with VSCSI, you must log into the host as `padmin` and then enter the `oem_setup_env` command to become root.

2. Go to the directory containing the software package you want to install. This will be either the MPIO, SAN, or SAN_Tool_Kit directory that was created when you uncompressed the file containing the Host Utilities software packages. These directories are subdirectories of the `ibm_aix_host_utilities_5.1` directory, which is located in the directory where you uncompressed the download file.

Below are command lines for installing the different software packages.

(AIX MPIO) To install the host settings software for an AIX MPIO environment, enter the following command line:

```
installp -aXYd /var/tmp/ibm_aix_host_utilities_5.1/MPIO/
Ontap.MPIO_Host_Utils_Kit
```

(Veritas DMP) To install the host settings software for a Veritas environment, enter the following command line:

```
installp -aXYd /var/tmp/ibm_aix_host_utilities_5.1/SAN /
Ontap.FCP_Host_Utils_Kitt
```

(Single-path iSCSI) To install the host settings software for a single-path iSCSI environment, enter the following command line:

```
installp -aXYd /var/tmp/ibm_aix_host_utilities_5.1/SAN /
Ontap.iSCSI_Host_Utils_Kit
```

(All environments) To install the SAN Toolkit for all Host Utilities environments, enter the following command line:

```
installp -aXYd /var/tmp/ibm_aix_host_utilities_5.1/SAN_Tool_Kit/
Ontap.SAN_toolkit
```

Example

The following is an example of installing the SAN Toolkit using the command line.

```
# installp -aXd Ontap.SAN_toolkit Ontap.SAN_toolkit
+-----+
          Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCCESSES
-----
Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets
-----
Ontap.SAN_toolkit.sanlun 5.1.0.0      #  SAN Toolki...
Ontap.SAN_toolkit.scripts 5.1.0.0      #  SAN Toolki...
** End of Success Section **

FILESET STATISTICS
```

```

-----+
 2 Selected to be installed, of which:
 2 Passed pre-installation verification
-----
 2 Total to be installed

+-----+
           Installing Software...
+-----+

installp: APPLYING software for:
  Ontap.SAN_toolkit.scripts 5.1.0.0
  Ontap.SAN_toolkit.sanlun 5.1.0.0
. . . . . ** Copyright notice for Ontap.SAN_toolkit ** . . . . .
(C) Copyright NetApp, Inc. 2003-2011.
All rights reserved.
. . . . . ** End of copyright notice for Ontap.SAN_toolkit **. . .
Finished processing all filesets. (Total time: 2 secs).
+-----+
           Summaries:
+-----+
Installation Summary
-----

```

Name	Level	Part	Event	Result
Ontap.SAN_toolkit.scripts	5.1.0.0	USR	APPLY	SUCCESS
Ontap.SAN_toolkit.sanlun	5.1.0.0	USR	APPLY	SUCCESS

Values for the `queue_depth` and `num_cmd_elems` parameters

It is a good practice to confirm the values of `num_cmd_elems` and `queue_depth` parameters used with the HBA initiators.

The Host Utilities do not change the `num_cmd_elems` parameter value from the AIX default value.

For the `queue_depth` parameter, though, the host settings package provided by the Host Utilities changes the value to 64 on the hdisk handles.

You can use the command `lsattr -El fcsX` to verify these values.

The following example shows the value of `num_cmd_elems`.

```
# lsattr -El fcs0 -a num_cmd_elems
num_cmd_elems 200 Maximum number of COMMANDS to queue to the
adapter True
```

This example shows the queue depth on the hdisk handles.

```
# lsattr -El hdisk4 -a
queue_depth 64 Queue DEPTH True
```

(FC) AIX fc_err_recov parameter settings

The AIX Fast I/O Failure setting is controlled by the `fc_err_recov` attribute of each `fscsi` device handle. Setting `fc_err_recov` to a value of `fast_fail` is required by certain applications and MPIO environments, including Veritas DMP. Enabling `fast_fail` may decrease the I/O recovery time when a link loss occurs between the host and storage device.

IBM documents the `fc_err_recov fast_fail` setting in APAR IY85526, located at <https://www-304.ibm.com/support/docview.wss?uid=isgIIY85526>.

For more information on Symantec's requirements regarding the AIX Fast I/O Failure setting, see the Symantec article HOWTO21911, which is titled *Configuring the AIX fast fail feature for use with Veritas Volume Manager (VxVM) and Dynamic Multipathing (DMP)*.

In this guide, the section *(Veritas DMP) Installing the ASL and APM software* contains information on setting `fc_err_recov` to `fast_fail`.

Related tasks

[\(Veritas DMP\) Installing the ASL and APM software](#) on page 55

Information on removing the AIX Host Utilities

You can easily remove an older version of the Host Utilities using either SMIT or the `installp` command.

The Host Utilities are packaged as a standard AIX fileset. It is a good practice to remove the Host Utilities software before you install a new version of the Host Utilities or switch from one environment, such as Veritas, to another environment. By performing a clean upgrade, you can avoid possible problems in the future.

If you are using a SAN boot LUN, you must perform special steps when you want to upgrade to a new version of the Host Utilities. See the section *(MPIO) Updating a SAN boot LUN* for more information.

Next topics

[*Before you uninstall previous versions of software*](#) on page 37

[*Uninstalling the Host Utilities software using SMIT*](#) on page 37

[*Using a command line to uninstall the Host Utilities software*](#) on page 41

Before you uninstall previous versions of software

Before you uninstall previous versions of the Host Utilities, you should prepare your system.

You need to disconnect the Data ONTAP LUNs from the host before you uninstall the Host Utilities. To disconnect the LUNs, you must:

1. Stop all I/O on the system.
2. Unmount the file systems that contain the volume groups mapped to storage system LUNs.
3. Varyoff the volume groups.
4. Remove the existing Data ONTAP LUN handles.

Uninstalling the Host Utilities software using SMIT

You can use SMIT to uninstall the Host Utilities software for your environment.

About this task

The example used in these steps uninstalls the MPIO host settings software package. If you have a different environment, substitute the name of that environment's host settings software package. You must repeat the steps to uninstall the SAN Toolkit.

Steps

1. Log in to the host as root.

(PowerVM) In PowerVM environments with VSCSI, you must log into the host as `padmin` and then enter the `oem_setup_env` command to become root.

2. Start SMIT by entering the following command:

`smit`

3. At the screen that appears, select the option titled Software Installation and Maintenance.

Example

When you start SMIT, it displays the following screen. On this screen, the Software Installation and Maintenance option is the first menu option.

```
+-----+
# smit
      System Management

Move cursor to desired item and press Enter.
      Software Installation and Maintenance
      Software License Management
      Devices
      System Storage Management (Physical & Logical Storage)
      Security & Users
      Communications Applications and Services
      Print Spooling
      Problem Determination
      Performance & Resource Scheduling
      System Environments
      Processes & Subsystems
      Applications
      Cluster System Management
      Using SMIT (information only)

F1=Help      F2=Refresh      F3=Cancel
Esc+8=Image  Esc+9=Shell    Esc+0=Exit      Enter=Do
+-----+
```

4. At the screen that appears, select the Software Maintenance and Utilities menu option.

Example

The following is an example of the Software Installation and Maintenance screen.

```
+-----+
Software Installation and Maintenance
Move cursor to desired item and press Enter.

      Install and Update Software
      List Software and Related Information
      Software Maintenance and Utilities
      Network Installation Management
      System Backup Manager
+-----+
```

```

F1=Help      F2=Refresh      F3=Cancel
Esc+8=Image  Esc+9=Shell      Esc+0=Exit      Enter=Do
+-----+

```

- At the next screen, select the Remove Installed Software menu option.

Example

The following is an example of the Software Maintenance and Utilities screen.

- Remove the software by entering the package name.

For example, to remove the mpio software package, you might enter

Ontap.MPIO_Host_Utils

Make sure that the Preview only option is set to no.

Example

(MPIO) The following is an example of the Remove Installed Software screen that appears when you are removing the Host Utilities software package for environments using MPIO for multipathing.

Note: By default PREVIEW only? is set to yes. You must change it to no if you want to uninstall the software.

```

+-----+
Remove Installed Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                         [Entry Fields]
* SOFTWARE name
[Ontap..MPIO_Host_Utils
.config Ontap.mpio_Host_Utils_kit.fcp
Ontap.MPIO_Host_Utils
_>

+
PREVIEW only? (remove operation will NOT occur)
no   +
REMOVE dependent software?
no   +
EXTEND file systems if space needed?
no   +
DETAILED output?
no   +

```

```

F1=Help      F2=Refresh      F3=Cancel      F4=List
Esc+5=Reset  Esc+6=Command  Esc+7=Edit
Esc+8=Image

```

Esc+9=Shell Esc+0=Exit Enter=Do
+-----+

Example of successful uninstall: The following output appears when you successfully uninstall the software for an MPIO environment.

```
Command: running      stdout: yes      stderr: no

Before command completion, additional instructions may appear below.

0503-409 installlp: bosboot verification starting...
installlp: bosboot verification completed.
+-----+ Deinstalling Software...
+-----+
Ontap.MPIO_Host_Utils_Kit.config 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.fcp 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.iscsi 5.1.0.0
Ontap.MPIO_Host_Utils_Kit.pcmodm 5.1.0.0

Finished processing all filesets. (Total time: 4 secs).

0503-409 installlp: bosboot verification starting...
installlp: bosboot verification completed.
0503-408 installlp: bosboot process starting...
bosboot: Boot image is 23825 512 byte blocks.
0503-292 This update will not fully take effect until after a
system reboot.

* * * A T T E N T I O N * * *
System boot image has been updated. You should reboot the
system as soon as possible to properly integrate the changes
and to avoid disruption of current functionality.

installlp: bosboot process completed.
+-----+
Summaries:
+-----+
Installation Summary
+-----+
Name      Level      Part      Event
Result
+-----+
Ontap.MPIO_Host_Utils_Kit.config 5.1.0.0 USR DEINSTALL SUCCESS
Ontap.MPIO_Host_Utils_Kit.fcp 5.1.0.0 USR DEINSTALL SUCCESS
Ontap.MPIO_Host_Utils_Kit.iscsi 5.1.0.0 USR DEINSTALL SUCCESS
Ontap.MPIO_Host_Utils_Kit.pcmod 5.1.0.0 USR DEINSTALL SUCCESS

F1=Help      F2=Refresh      F3=Cancel      Esc
+6=Command
Esx Esc+8=Image Esc+9=Shell      Esc+0=Exit
```

Using a command line to uninstall the Host Utilities software

You can use a command line to uninstall the Host Utilities software for your environment.

Steps

1. Log in to the host as root.

(PowerVM) In PowerVM environments with VSCSI, you must log into the host as `padmin` and then enter the `oem_setup_env` command to become root.

2. Use the `installp -u FileSetName` command to uninstall the software package.

Below are the command lines for uninstalling the different software packages.

(MPIO) To uninstall the host settings software for an MPIO environment, enter the following command line:

```
installp -u Ontap.MPIO_Host_Utils
```

(Veritas) To uninstall the host settings software for a Veritas environment, enter the following command line:

```
installp -u Ontap.FCP_Host_Utils
```

(Single-path iSCSI) To uninstall the host settings software for a single-path iSCSI environment, enter the following command line:

```
installp -u Ontap.iSCSI_Host_Utils
```

(All environments) To uninstall the SAN Toolkit for all Host Utilities environments, enter the following command line:

```
installp -u Ontap.SAN_toolkit
```


(iSCSI) Additional configuration for iSCSI environments

When you are using the iSCSI protocol, you need to perform some additional tasks to complete the installation of the Host Utilities. Some of the tasks you perform vary depending on whether you have a software iSCSI initiator or a hardware iSCSI HBA.

To complete the setup of the Host Utilities in an iSCSI environment, you must:

- **(Software iSCSI)** Record the host's initiator node name. You need this information to set up your storage.
- **(Software iSCSI)** Change the software initiator's iSCSI node name from the default node name.
- **(Hardware iSCSI HBA)** Prepare the iSCSI HBA by either removing the protocol device if it is configured or changing the database.
- **(Hardware iSCSI HBA)** Configure the iSCSI HBA.
- **(Software and Hardware iSCSI)** Add information to the targets file.
- **(Software and Hardware iSCSI)** (Optional) Configure CHAP.

The following sections explain how to perform these tasks.

Next topics

[*\(iSCSI\) Changing the software initiator's iSCSI node name*](#) on page 43

[*\(iSCSI\) Prepare the HBA for configuration*](#) on page 44

[*\(iSCSI\) Configuring the HBA*](#) on page 45

[*\(iSCSI\) Adding information to the targets file*](#) on page 47

[*\(iSCSI\) CHAP authentication*](#) on page 49

(iSCSI) Changing the software initiator's iSCSI node name

You must change the default initiator node name. The default node name generated by AIX does not work with the storage system.

About this task

Confirm that your system is correctly set up for iSCSI by making sure the node name is correct.

Steps

1. From the AIX host console, enter:

```
lsdev -c | grep -i iscsi
```

Example

The following is an example of the type of information you should see when you execute these commands.

```
AIXHOST >lsdev -C |grep -i iscsi
iscsi0 Available iSCSI Protocol Device
```

2. Get the current node name. Enter:

```
lsattr -El iscsi0
```

By having the current node name to compare against the new name, you can verify that the name change took place.

Example

The following is an example of the type of information you should see when you execute this command.

```
AIXHOST >lsattr -El iscsi0
disc_filename /etc/iscsi/targets Configuration file False
disc_policy file Discovery Policy True
initiator_name iqn.aixp52.hostid.0a3da60c iSCSI Initiator Name True
max_targets 16 Maximum Targets Allowed True
num_cmd_elems 200 Maximum number of commands to queue to driver True
```

3. Change the node name. Enter:

```
chdev -l iscsi0 -a initiator_name=nodename
```

nodename is the new iSCSI node name for the host. Node names must follow these rules:

- The format must be iqn.1996-04.com.ibm:*unique-value*.
- The maximum number of characters is 223.
- Allowed characters are lowercase letters (a to z), numbers (0 to 9), period (“.”), hyphen (“-”), and colon (“:”).
- The underscore character (“_”) is not allowed.

4. Verify the new node name by displaying it and comparing it to the original name. Enter:

```
lsattr -El iscsi0
```

(iSCSI) Prepare the HBA for configuration

Before you can configure the HBA, you must make sure it is in the right state.

You must do one of the following:

- Remove the protocol device if it is configured. Normally, this device will be configured. To find the protocol device for the HBA, use the `lsdev -p <adapter>` command. To remove the protocol device, use the `rmdev -dl <protocol_device>` command. Once you've found and removed the protocol device, you can configure the adapter without any issues.

- If the protocol device is not configured, make changes to the database. Once you make the changes, you must reboot your system.

Example of locating and removing a protocol device

The following command line locates the protocol device iscsi0.

```
# lsdev -p ics0
iscsi0 Available 01-11-01 iSCSI Protocol Device
```

The next command removes the protocol device iscsi0.

```
# rmdev -dl iscsi0
iscsi0 Deleted
```

(iSCSI) Configuring the HBA

After you insert the HBA and install the Host Utilities software, you must configure the adapter.

Step

1. Use the smit iscsi command to configure the HBA.

Example

The following example shows you how to select the option to configure the characteristics of the adapter and select the number of the adapter you are configuring. The options you need to select are shown in **bold**. Also, the information that you must enter once you select "Change / Show Characteristics of an iSCSI Adapter" is shown in **bold**.

Note: This is an example of configuring the adapter after you have removed the protocol device. If you did not remove the protocol device, then you would set "Apply change to DATABASE only" to yes.

```
# smit iscsi
+-----+
          iSCSI
Move cursor to desired item and press Enter.

  iSCSI Adapter
  iSCSI Protocol Device
  iSCSI Target Device Parameters in ODM
+-----+
Select "iSCSI Adapter"
+-----+
          iSCSI Adapter
Move cursor to desired item and press Enter.
```

```

List All iSCSI Adapters
Change / Show Characteristics of an iSCSI Adapter
Generate Error Report
Trace iSCSI Adapter
+-----+
+-----+
|                                         iSCSI Adapter
|                                         +-----+
|                                         Move cursor to desired item and press Enter.

|                                         ics0 Available 01-11 iSCSI Adapter
|                                         ics1 Available 01-13 iSCSI Adapter
|                                         +-----+
|                                         F1=Help          F2=Refresh        F3=Cancel
|                                         F8=Image          F10=Exit          Enter=Do
|                                         /=Find           n=Find Next
|                                         +-----+
|                                         Change / Show Characteristics of an iSCSI
|                                         Adapter
|                                         +-----+
|                                         Type or select values in entry fields.
|                                         Press Enter AFTER making all desired changes.

                                         [Entry Fields]
iSCSI Adapter                               ics0
Description                                 iSCSI Adapter
Status                                     Available
Location 09-09
iSCSI Initiator Name                      [iqn.1996-04.com.ibm:sn.
0310b806g1]
                                         Maximum number of Commands to Queue to Adapter
                                         [200]          +#
                                         Maximum Transfer Size
                                         [0x100000]      +
                                         Discovery Filename
                                         Discovery Policy
file                                     +
                                         Automatic Discovery Secrets Filename
                                         Adapter IP Address
                                         Adapter Subnet Mask
                                         Adapter Gateway Address
                                         [ ]
                                         [10.60.244.136]
                                         [255.255.255.0]
                                         [ ]

```

(iSCSI) Adding information to the targets file

You must edit the targets file indicated by the Discovery Filename parameter to provide information on the IP address, TCP port, and iSCSI target node name for each storage system in the cluster pair.

About this task

Note: The iSCSI target only supports file as the value of the Discovery Policy.

Steps

1. Get the IP address and the TCP port. You can use the `iscsi portal show` command on the storage system to get this information.

Example

The following example uses the `iscsi portal show` command to display information about the portal.

```
# iscsi portal show
  Network portals:
    IP address      TCP Port      TPGroup      Interface
      10.60.244.34      3260        1000        e0
```

2. Get the node name. You can use the `iscsi nodename` command on the storage system or the `lsdev` and `grep` commands on the host to do this.

Example

The following example uses the `lsdev` and `grep` commands to display information about the iSCSI nodes.

```
# lsdev | grep 'iSCSI Protocol Device'
iscsi0      Available 11-09-01      iSCSI Protocol Device
iscsi1      Available 1D-09-01      iSCSI Protocol Device
iscsi2      Available          iSCSI Protocol Device (software
initiator>
```

3. Check the final `disk_filename` file. You can use the `lsattr -El` command on the host.

Example

The following example checks the `final disk_filename` file.

```
# lsattr -El iscsi2 -a disk_filename
disk_filename /etc/iscsi/targets Configuration file False
```

Examples of checking software initiators and HBAs

HBA: The following is an example of checking the IP address, TCP port, node name, and `disk_filename` file when you are using an HBA.

If your environment uses multiple HBAs, you may need to set up multiple target files. This is only necessary if you have multiple unique network segments. In that case you would need a

unique target file for each network segment. On a flat network, you would want one targets file and all the adapters to log into the same iSCSI target.

To log back into the target, you must rescan the bus using the `cfgmgr` command.

```
# lsdev -Cc adapter | grep 'iSCSI Adapter'
ics0    Available 11-09      iSCSI Adapter
ics1    Available 1D-09      iSCSI Adapter

# lsattr -El ics0 -a disc_filename
disc_filename /etc/iscsi/targetshw_VLAN_100 Discovery Filename
True

# lsattr -El ics1 -a disc_filename
disc_filename /etc/iscsi/targetshw_VLAN_200 Discovery Filename
True

# cat /etc/iscsi/targetshw_VLAN_100
192.168.100.222 3260 iqn.1986-03.com.ibm:sn.84209003
192.168.100.221 3260 iqn.1986-03.com.ibm:sn.84213404

# cat /etc/iscsi/targetshw_VLAN_200
192.168.200.222 3260 iqn.1986-03.com.ibm:sn.84209003
192.168.200.221 3260 iqn.1986-03.com.ibm:sn.84213404
```

(iSCSI) Adding targets

Before you can access LUNs on a storage system, you must add an entry for the storage system to the Discovery Filename on the AIX host.

About this task

You need one entry for each storage system, regardless of the number of interfaces that are enabled for iSCSI traffic. A clustered storage system needs two entries, one for each node in the cluster.

Note: For more information on working with targets for the HBA, see the “targets File” section of the *IBM AIX Files Reference*.

Steps

1. With a text editor, open the file specified by the `disk_filename` parameter.
2. Add one line for one interface on each storage system. Be sure to use an interface that is enabled for iSCSI traffic. The line has the following format:

`HostNameOrAddress PortNumber iSCSIName ["CHAPSecret"]`

`HostNameOrAddress` is the host name or IP address of a gigabit Ethernet interface on the storage system. Specify an interface that is enabled for iSCSI communication.

PortNumber is always 3260.

iSCSIName is the iSCSI target node name of the storage system.

CHAPSecret is the optional CHAP password for the host. Enclose the text string in double quotes. This value must match the value configured on the storage system for this initiator.

(iSCSI) CHAP authentication

If you choose, you can also configure CHAP authentication. The AIX initiator supports both unidirectional and bidirectional CHAP.

You must configure CHAP on both the host and the storage system. The direction type you use indicates the direction of authentication relative to the storage system.

While the storage system recognizes two types of CHAP user names and passwords, you must specify inbound if you are using CHAP authentication. With inbound, the storage system authenticates the initiator or HBA.

Note: Do not specify outbound settings for AIX hosts. Neither the AIX software initiator nor the HBA supports authentication of the storage system using CHAP.

The initiator CHAP secret value you configure on the host must be the same as the inpassword value you configured on the storage system. The initiator CHAP name must be the same as the inname value you configured on the storage system.

Note: The iSCSI initiator allows a single CHAP secret value that is used for all targets. If you try to configure a second CHAP secret, it overwrites the first value you set.

CHAP settings are defined in the `/etc/iscsi/targets` file on the host. This file is specified in the Discovery Filename parameter. The storage system's inbound password must match the initiator's CHAPSecret in the `/etc/iscsi/targets` file on the host. The storage system's inbound user name must match the initiator node name of the host. The AIX initiator or HBA always uses its iSCSI node name as its CHAP user name.

(iSCSI) Data ONTAP upgrades may affect CHAP configuration

In some cases, if you upgrade the Data ONTAP software running on the storage system, the CHAP configuration on the storage system is not saved.

To avoid losing your CHAP settings, run the `iscsi security add` command. You should do this even if you have already configured the CHAP settings.

(Veritas DMP) Configuration requirements for Veritas Storage Foundation environments

There are several tasks you must perform to set up your Veritas DMP environment. Some of them, such as whether you need to install the Array Support Library (ASL) and the Array Policy Module (APM), depend on your version of Veritas Storage Foundation. Others, such as setting the `fc_err_recov` parameter to `fast_fail`, are required for all versions of Veritas Storage Foundation.

To determine whether you need to install the ASL and APM, check your version of Veritas Storage Foundation:

- If you have Veritas Storage Foundation 5.1, you do not need to install the ASL and APM. They are included with the Veritas Storage Foundation product.
- If you have Veritas Storage Foundation 5.0, you must manually install the ASL and APM.

With the ASL and APM installed, you can use either the `sanlun` utility or VxVM to display information about the paths to the LUNs on the storage system.

In addition to confirming that you have the correct ASL and APM installed for your system, you should also set the Veritas restore daemon values for the restore policy and the polling interval to the Host Utilities recommended values. The section *(Veritas DMP) Setting the restore daemon interval* contains information the values you should use.

For information on setting the `fc_err_recov` parameter to `fast_fail`, see the section *(Veritas DMP) Installing the ASL and APM software*. Setting `fc_err_recov` is one of the steps.

Next topics

(Veritas DMP) The Array Support Library and the Array Policy Module on page 51

(Veritas DMP) Using sanlun to obtain multipathing information on page 58

(Veritas DMP) Using VxVM to display available paths on page 58

(Veritas DMP) The Veritas DMP restore daemon requirements on page 59

(Veritas DMP) Information on ASL error messages on page 60

(Veritas DMP) The Array Support Library and the Array Policy Module

The ASL and APM for N series storage systems are necessary if you want to use Veritas with the Host Utilities. While the ASL and APM are qualified for the Host Utilities, they are provided and supported by Symantec.

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

Note: If you encounter a problem with the ASL or APM, contact Symantec customer support.

To determine which versions of the ASL and APM you need for your version of the Host Utilities, check the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html. This information is updated frequently. Once you know the version you need, go to the Symantec Web site and download the ASL and APM.

Note: With Veritas Storage Foundation 5.0, you must install both the ASL and the APM.

The ASL is an IBM-qualified library that provides information about storage array attributes configurations 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. When the ASL detects the N series storage array, it sets the array type to A/A-NETAPP or ALUA. This instructs Veritas software to use the N series APM to handle I/O and path failures for the storage array.

Next topics

[\(Veritas DMP\) Information provided by the ASL](#) on page 52

[\(Veritas DMP\) Information on installing and upgrading the ASL and APM](#) on page 53

[\(Veritas DMP\) Removing the ASL and APM](#) on page 57

Related tasks

[\(Veritas DMP\) Installing the ASL and APM software](#) on page 55

Related references

[\(Veritas DMP\) The Veritas DMP restore daemon requirements](#) on page 59

(Veritas DMP) Information provided by the ASL

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

Note: You cannot use storage systems simultaneously as JBODs and vendor arrays. If you install the ASL, storage systems are reported as storage arrays unless you explicitly exclude them by using the `vxddladm excludearray` command. You cannot configure them in VxVM as JBODs.

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.

- Array Type.

- Active/active (A/A-NETAPP)

All optimized (primary or active) paths can be used for I/O concurrently. When all optimized paths to a LUN in a clustered configuration are unavailable, the un-optimized (secondary or standby) paths are used.

Note: In addition to the clustered configurations, the ASL supports direct-attached, non-clustered configurations, including NearStore models. These configurations are also reported as active/active array types.

- ALUA

A LUN in an ALUA-enabled array can be accessed through both controllers, using optimized and non-optimized paths. The array informs the host of path options, their current state, and state changes. Using this information, the host can determine which paths are optimized.

Failover to the non-optimized path occurs only if all the optimized paths fail.

For 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) Information on installing and upgrading the ASL and APM

If you are using a Veritas environment, you must use the ASL and APM. While the ASL and APM are included with Veritas Storage Foundation 5.1, other versions of Veritas Storage Foundation require that you install them.

If you are using Veritas Storage Foundation 5.0, you must install both the ASL and the APM.

Before you can install the ASL and APM, you must first remove any currently installed versions of 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 NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html for current information about the system requirements.
- If you currently have the ASL installed, determine its version to see if it is the most up-to-date version for your system.
- 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.

- Get the new ASL and the APM from Symantec.
- Follow the instructions in the Symantec TechNote as well as the steps provided in this chapter to install the new version of the ASL and APM.

Next topics

- [*\(Veritas DMP\) Determining the ASL version* on page 54](#)
- [*\(Veritas DMP\) How to get the ASL and APM* on page 54](#)
- [*\(Veritas DMP\) Installing the ASL and APM software* on page 55](#)

(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
=====
libvxshark.so    vm-5.0-rev-1    5.0
libvxcscovrts.so vm-5.0-rev-1    5.0
libvxemc.so      vm-5.0-rev-2    5.0
libvxhitachi.so  vm-5.0-rev-1    5.0
libvxibmds4k.so  vm-5.0-rev-1    5.0
libvxibmds6k.so  vm-5.0-rev-1    5.0
libvxibmds8k.so  vm-5.0-rev-1    5.0
libvxhds9980.so  vm-5.0-rev-1    5.0
libvxhdsalua.so  vm-5.0-rev-1    5.0
libvxhdsusp.so   vm-5.0-rev-2    5.0
libvxhds.so      vm-5.0-rev-1    5.0
libvxCLARIOn.so  vm-5.0-rev-1    5.0
libvxpurple.so   vm-5.0-rev-1    5.0
libvxhdssusp.so  vm-5.0-rev-1    5.0
libvxxp256.so    vm-5.0-rev-1    5.0
libvxibmsvc.so   vm-5.0-rev-1    5.0
libvxvscsi.so    vm-5.0-rev-1    5.0
libvxvpath.so    vm-5.0-rev-1    5.0
libvxnetapp.so   vm-5.0-rev-0    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.

Note: If you have Veritas Storage Foundation 5.1, you do not need to install the ASL and APM. They are installed with the Veritas Storage Foundation software.

Information on getting the Symantec TechNote for the ASL and APM TechNote is provided on the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.

(Veritas DMP) Installing the ASL and APM software

If you do not have the correct version of the ASL and APM, you must install them. The installation involves several steps.

Before you begin

Make sure you have downloaded the correct version of the ASL and APM. To confirm which version of the ASL and APM you need, check the NAS Interoperability Matrices Web site.

You also need the Symantec ASL/APM TechNote, which you can view at the Symantec Web site. The TechNote contains the Symantec instructions for installing the ASL and APM.

Note: You only need to manually install the ASL and APM if you are using Veritas Storage Foundation 5.0 or if you need to re-install them. With Veritas Storage Foundation 5.1, the ASL and APM are automatically installed when you install that product.

Steps

1. Log in to the VxVM system as the root user.
2. If you have VxVM installed, stop all applications using VxVM volumes. Make sure that there are no VxVM volumes open, mounted, or in use.

Note: Check the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html to make sure you have the supported maintenance packages. You should install the correct VxVM and maintenance packages before you install the ASL and APM.

3. For each host HBA (fscsiX), enable the AIX Fast Fail feature by setting the `fc_err_recov` parameter to `fast_fail`.

Note: When you run Veritas Storage foundation, you must set the `fc_err_recov` parameter to `fast_fail`. This feature improves I/O failover handling.

Example

The following example sets the `fc_err_recov` parameter. The `-P` flag on the command line changes the ODM stanza, but not the active configuration. You must reboot your system for this change to take place.

```
chdev -l fscsi0 -a fc_err_recov=fast_fail -P
chdev -l fscsi1 -a fc_err_recov=fast_fail -P
```

4. For each host HBA (fscsiX), enable Dynamic Tracking.

Example

The following example enables Dynamic Tracking.

```
chdev -l fscsi0 -a dyntrk=yes -P
chdev -l fscsi1 -a dyntrk=yes -P
```

5. Install the ASL and APM according to the installation instructions provided by the ASL/APM TechNote on the Symantec Web site.

6. Reboot the host.

Note: If you do not reboot the host, the changes will not take effect.

7. Verify that FastFail and Dynamic Tracking are enabled by entering the following command for each HBA:

```
lsattr -El fscsiX
```

X is the controller/HBA ID.

Example

The following output indicates that Dynamic Tracking is on (dyntrk value=yes) and the FastFail feature is enabled (fc_err_recov=fast_fail).

```
attach      switch      How this adapter is CONNECTED      False
dyntrk     yes         Dynamic Tracking of FC Devices  True
fc_err_recov fast_fail  FC Fabric Event Error RECOVERY Policy  True
scsi_id    0x10400    Adapter SCSI ID                  False
sw_fc_class 3          FC Class for Fabric          True
```

8. Verify that the ASL is installed by entering the following command:

```
vxddladm listsupport
```

Example

As shown in the output below, the libvxnetapp.so library is part of the command output.

```
## vxddladm listsupport
LIBNAME          VID
=====
libvxshark.so   IBM
libvxcscovrts.so CSCOVRTS
libvxhitachi.so HITACHI
libvxbmds6k.so  IBM
libvxp256.so    IBM
libvxp256.so    HP
libvxbmds4k.so  IBM
libvxemc.so     EMC
libvhdsuspl.so  HITACHI
libvhds.so      HITACHI
libvhdsalua.so  HITACHI
libvhds9980.so  HITACHI
libvxbmds8k.so  IBM
```

libvxpurple.so	SUN
libvxCLARiON.so	DGC
libvxhpalua.so	HP, COMPAQ
libvxibmsvc.so	IBM
libvxvscsi.so	AIX
libvxnetapp.so	NETAPP

- Verify that the APM is installed by entering the following command:

```
vxdmpadm listapm all
```

Example

When you execute the `vxdmpadm listapm all` command, it displays information on the array type and the state of the APM. If the storage controller disks have not been detected by the host, the state of the APM is Not-Active. As shown in the output that follows, there is also an Array Type of A/A-NETAPP.

Module Name	APM Name	APM Version	Array Types	State
dmppaa	dmppaa	1	A/A	Active
dmppap	dmppap	1	A/P	Active
dmppap	dmppap	1	A/P-C	Active
dmppapf	dmppapf	1	A/PF-VERITAS	Not-Active
dmppapf	dmppapf	1	A/PF-T3PLUS	Not-Active
dmppapg	dmppapg	1	A/PG	Not-Active
dmppapg	dmppapg	1	A/PG-C	Not-Active
dmpjbod	dmpjbod	1	Disk	Active
dmpjbod	dmpjbod	1	APdisk	Active
dmpnetapp	dmpnetapp	1	A/A-NETAPP	Not-Active
dmppsvc	dmppsvc	1	A/A-IBMSVC	Not-Active

(Veritas DMP) Removing the ASL and APM

The host can have only one version of an ASL that supports storage systems running Data ONTAP installed on it at a time. If you need to upgrade the ASL and APM, you must first remove them. This task involves several steps.

Before you begin

You do not need to stop any volumes on unaffected disk arrays, such as disk arrays from other vendors, before you remove the ASL and APM. This is also true for arrays or disks in the OTHER_DISKS category. OTHER_DISKS are local, non-FC-attached disks. Volumes created on these arrays remain accessible because they do not have multiple paths.

Steps

- Log in to the VxVM system as the root user.
- Stop all I/O to the N series LUNs.

Note: If you have a Veritas Storage Foundation RAC cluster, you must also stop clustering on the node before remove the ASL.

3. Use either SMIT or the command line to remove the ASL and APM packages.

(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 A/A-NetApp or ALUA.

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) 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:

```
vxldmpadm 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:

```
vxldmpadm getsubpaths ctrlr=controller_name
```

controller_name is the controller displayed under CTRLR-NAME in the output of the vxldmpadm 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) 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. By default, the restore daemon checks for disabled paths every 300 seconds.

The Host Utilities recommended settings for these values at the time this document was produced were a restore policy of "disabled" and a polling interval of "60".

Check the Release Notes to see if these recommendations have changed.

(Veritas DMP) Setting the restore daemon interval

You can change the value of the restore daemon interval to match the recommendation for the Host Utilities. Doing this improves the I/O failover handling.

About this task

At the time this document was prepared, it was recommended that you set the restore daemon interval value to 60 seconds to improve the recovery of previously failed paths and the restore policy to disabled. The following steps take you through the process of setting the values.

Steps

1. Stop the restore daemon.

```
/usr/sbin/vxldmpadm stop restore
```

2. Change the restore daemon setting to 60 and set the policy to `check_disabled`.

```
/usr/sbin/vxldmpadm start restore interval=60 policy=check_disabled
```

Note: This step reconfigures and restarts the restore daemon without the need for an immediate reboot.

3. Make sure the new restore daemon interval will persist across reboots. The method you use to do this varies depending on your Veritas Storage Foundation configuration.

Option	Description
Veritas Storage Foundation 5.0 MP1 and earlier	<p>Perform these steps to ensure that new restore daemon interval persists across reboots.</p> <ol style="list-style-type: none"> <li data-bbox="395 273 1233 307">Edit the <code>/etc/init.d/vxvm-sysboot</code> file to set the interval to 60 seconds. <pre data-bbox="435 307 1233 331">restore_daemon_opts="interval=60 policy=check_disabled"</pre> <p>By default, the restore daemon options stored in this file are:</p> <pre data-bbox="435 376 915 432">restore_daemon_opts="interval=300 policy=check_disabled"</pre> <ol style="list-style-type: none"> <li data-bbox="395 451 995 468">Save and exit the <code>/etc/init.d/vxvm-sysboot</code> file.
Veritas Storage Foundation 5.0 MP3 and 5.1	<p>Enter the following command line to change the restore daemon interval to 60 seconds:</p> <pre data-bbox="395 584 973 608">vxldmpadm settune dmp_restore_interval=60</pre> <p>When you use the <code>vxldmpadm</code> command to make the change with this version of Veritas Storage Foundation, the change persists across reboots.</p>

4. Verify the changes.

```
/usr/sbin/vxldmpadm stat restored
```

The command output shows the status of the restore daemon. Below is a sample of the type of output the command displays.

```
The number of daemons running : 1
The interval of daemon: 60
The policy of daemon: check_disabled
```

(Veritas DMP) Information on 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.

Message severity	Definition
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.

Overview of LUN configuration and management

Configuring and managing LUNs involves several tasks. Whether you are executing the Host Utilities in a Veritas DMP environment or an MPIO environment determines which tasks you need to perform.

The following table summarizes the tasks for all the supported AIX Host Utilities environments. If a task does not apply to all environments, the table specifies which environments it does apply to. If a task does not apply to your Host Utilities environment, do not perform it.

Task	Discussion
1. Create and map igroups and LUNs	<p>An igroup is a collection of WWPNs on the storage system that map to one or more host HBAs.</p> <p>You must create at least one igroup and at least one LUN. Then you must map the LUN to the igroup.</p> <p>For complete information on creating igroups and LUNs, see the <i>Data ONTAP Block Access Management Guide for iSCSI and FC</i> for your version of Data ONTAP.</p>
2. (FC) Enable ALUA	<p>If your environment uses MPIO for multipathing and supports ALUA, you should set it up to work with igroups. With Data ONTAP 8.0 or later, ALUA is enabled by default when you create an igroup that works with the FC protocol. Earlier versions of Data ONTAP require you to manually enable ALUA. In Veritas environments, you can use Veritas DMP instead of ALUA.</p> <p>To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.</p> <p>Note: With certain versions of Data ONTAP, ALUA is automatically enabled when you create the igroup. You can use the command <code>igroup show -v <i>igroup_name</i></code> to see if ALUA is enabled. If it is not enabled, you can use the command <code>igroup set <i>igroup</i> alua yes</code> to enable it.</p>

Task	Discussion
3. Discover and configure LUNs	<p>You can discover LUNs using the following commands:</p> <ul style="list-style-type: none"> • Veritas: cfgmgr followed by vxdctl enable • PowerVM cfgdev • MPIO cfgmgr
4. (Optional) Create volume groups and file systems	You can create volume groups and file systems using the Veritas or AIX tools that you normally use for these tasks.
5. Display information about the LUNs and the HBAs	You can use the <code>sanlun</code> utility to display information about the LUNs and the HBAs.

For more information about what is supported, see the NetApp Interoperability Matrix at www.ibm.com/systems/storage/network/interphome.html.

Next topics

(AIX MPIO, PowerVM) Discovering and Managing LUNs on page 64

(Veritas DMP) Discovering and configuring LUNs with Veritas on page 72

(AIX MPIO, PowerVM) Discovering and Managing LUNs

There are several tasks you need to perform when setting up LUNs for an AIX MPIO or PowerVM environment. The following sections provide information on those tasks.

Next topics

(AIX MPIO without PowerVM) Discovering and configuring LUNs with AIX on page 65

(PowerVM VSCSI) Discovering and configuring LUNs on page 66

(PowerVM NPIV) Discovering and configuring LUNs on page 70

(AIX MPIO without PowerVM) Discovering and configuring LUNs with AIX

When you are using AIX without PowerVM, you use the `cfgmgr` command to get the host to discover the LUNs. You do not need to reboot the host to discover the LUNs. To confirm that the host discovered the LUNs, you can use the `sanlun` utility.

Before you begin

(iSCSI) Make sure you have added an entry for the storage system in the file specified by the `disk_filename` parameter.

About this task

The steps that follow do not apply to a host that is running a PowerVM environment.

Steps

1. Log in as root on the host.
2. On the AIX host, enter the following command to discover the newly created LUNs:

`cfgmgr`

The `cfgmgr` command causes

- The initiators to log into the fabric.
- The system to check for new devices.
- The system to create new device entries.

3. To display all AIX disks, including the LUNs, enter the following command:

`lsdev -Cc disk`

The LUNs appear on the host as disks. You can use these disks the same way you use local disks on the host. The host settings ODM file that is installed as part of the Host Utilities causes the LUNs to be displayed as either FCP disks or iSCSI disks, depending on which protocol you are using.

Write down the hdisk instance numbers. You use them when you perform the path configuration in Step 5.

4. To get information about your setup, enter the following command:

`lsattr -El hdisk_name`

5. Order and prioritize the paths to the storage system LUNs.

- If you are running the FC protocol and a version of Data ONTAP that supports ALUA, you should have ALUA running on the storage system. ALUA automatically sets the paths.

Note: With certain versions of Data ONTAP, ALUA is automatically enabled when you create the igroup. You can use the command `igroup show -v igroup_name` to see if

ALUA is enabled on the storage system. If it is not enabled, you can use the command `igroup set igrup_name alua yes` to enable it.

- If you are not running ALUA, you must execute the `dotpaths` utility to set the path priorities. The `dotpaths` utility is installed when you install the Host Utilities. If you enter `dotpaths` without any options, it sets the priority for all Data ONTAP LUNs.

Note: Do not use the `dotpaths` utility if ALUA is enabled.

The command line for this utility has the following format:

`dotpaths [-hqv] [hdiskN ...]`

You can use the command line options to:

- Specify a set of disks for which you want to set the priority. When you enter a list of hdisk names, the utility sets the priority only for those disks.
- Display all the priorities. Enter the `-v` option to enable verbose mode.
- Query all or individual disk priorities. Enter the `-q` option to perform this task.

6. To verify that the host has discovered the LUNs, enter the following command:

`sanlun lun show`

(PowerVM VSCSI) Discovering and configuring LUNs

When you are using PowerVM VSCSI and VIO servers, you must use the `padmin` login and the commands appropriate for it to configure and discover LUNs.

Before you begin

(iSCSI) Make sure you have added an entry for the storage system in the file specified by the `disk_filename` parameter.

Steps

1. Log in to the host as `padmin`.
2. On the AIX host, enter the following command to discover the newly created LUNs:

`cfgdev`

The `cfgdev` command causes:

- The initiators to log into the fabric.
- The system to check for new devices.
- The system to create new device entries where the new LUNs are mapped as MPIO devices using the default PCM.

Example

On a VIO server, use the `cfgdev` command to discover the LUNs.

`vioserver$ cfgdev`

3. To verify that the hdisk MPIO devices were created, enter the following command:

```
lsdev -type disk
```

The LUNs appear on the host as disks. You can use these disks the same way you use local disks on the host.

Write down the hdisk instance numbers. You use them when you perform the path configuration in Step 6.

The `lsdev -type disk` command produces the following type of output.

```
vioserver$ lsdev -type disk NetApp
name      status      description
hdisk0    Available   16 Bit LVD SCSI Disk Drive
hdisk1    Available   16 Bit LVD SCSI Disk Drive
hdisk2    Available   MPIO NetApp FCP Default PCM Disk
hdisk3    Available   MPIO NetApp FCP Default PCM Disk
hdisk4    Available   MPIO NetApp FCP Default PCM Disk
hdisk5    Available   MPIO NetApp FCP Default PCM Disk
hdisk6    Available   MPIO NetApp FCP Default PCM Disk
```

The output for this example has been truncated.

4. To get information about your setup, enter the following command:

```
lsdev -dev hdisk_name -attr
```

The `lsdev -dev hdisk_name -attr` command produces the following output.

```
vioserver$ lsdev -dev hdisk2 -attr
attribute value          description          user-settable
PCM      PCM/friend/NetAppDefault  PCM Path Control Module  False
PR_key_value      none          Persistant Reserve Key
Value      True
algorithm      round_robin
Algorithm
clr_q      no          Device CLEARS its Queue on
error      True
dist_err_pcnt      0          Distributed Error Sample
Time      True
dist_tw_width      50         Distributed Error Sample
Time      True
hcheck_cmd      inquiry
Command
hcheck_interval 30
Interval
hcheck_mode      nonactive
Mode
location
Label
lun_id      0x2000000000000000
ID
lun_reset_spt      yes
Reset
max_transfer      0x100000
Size
node_name      0x500a0980894ae0e0
Name
pvid
identifier
q_err      yes
Use QERR
```

bit		True	
q_type	simple	True	Queuing
TYPE			
qfull_dly	2		Delay in seconds for SCSI TASK
SET FULL	True		
queue_depth	64		Queue
DEPTH		True	
reassign_to	120		REASSIGN time out
value	True		
reserve_policy	no_reserve		Reserve
Policy		True	
rw_timeout	30		READ/WRITE time out
value	True		
scsi_id	0xd10001		SCSI
ID		False	
start_timeout	60		START unit time out
value	True		
ww_name	0x500a0984994ae0e0		FC World Wide
Name		False	

5. Switch to the `oem_setup_env` mode so that you can execute the Host Utilities commands `dotpaths` (if you are not using ALUA) and `sanlun`. To get into this mode, enter the following command:

`oem_setup_env`

Note: Once you enter the `oem_setup_env` command, you stay in that mode until you halt it by entering a command such as `exit`.

Example

When you enter the `oem_setup_env` mode, the command line prompt changes from a dollar sign to a pound sign.

```
vioserver$ oem_setup_env
#
```

6. Order and prioritize the paths to the storage system LUNs.

- If you are running the FC protocol and a version of Data ONTAP that supports ALUA, you should have ALUA running on the storage system. ALUA automatically sets the paths.

Note: With Data ONTAP 8.0 and later, ALUA is automatically enabled when you create the `igroup`. To confirm that ALUA is enabled, run the command `igroup show -v igrup_name` on the storage system. If it is not enabled, you can use the command `igroup set igrup_name alua yes` to enable it.

- If you are not running ALUA, you must execute the `dotpaths` utility to set the path priorities. The `dotpaths` utility is installed when you install the Host Utilities. If you enter `dotpaths` without any options, it sets the priority for all Data ONTAP LUNs.

Note: Do not use the `dotpaths` utility if ALUA is enabled.

The command line for this utility has the following format:

`dotpaths [-hqv] [hdiskN ...]`

You can use the command line options to:

- Specify a set of disks for which you want to set the priority. When you enter a list of hdisk names, the utility sets the priority only for those disks.
- Display all the priorities. Enter the **-v** option to enable verbose mode.
- Query all or individual disk priorities. Enter the **-q** option to perform this task.

Example

The following is an example of using the **dotpaths** utility.

```
vioserver# dotpaths -v
hdisk2 (path 0): Optimized path - Setting priority to 255
hdisk2 (path 1): Optimized path - Setting priority to 255
hdisk2 (path 2): Unoptimized path - Current priority of 1 is correct
hdisk2 (path 3): Unoptimized path - Current priority of 1 is correct
...
Path priority set and/or verified for 22 disks, 86 total paths.
vioserver#
```

7. To verify that the host has discovered the LUNs, enter the following command:

```
sanlun lun show -p
```

Example

This example shows typical output of the **sanlun lun show -p** command for a direct-attached configuration. It also includes the **sanlun lun show -d hdisk9 -v** command to display information about hdisk9. The output from this command has been truncated.

```
# sanlun lun show -p
ONTAP_PATH: sh-960-01:/vol/vol1/vioserver_7
LUN: 7
LUN Size: 250m (262144000)
Host Device: hdisk9
LUN State: GOOD
Controller_CF_State: Cluster Enabled
Controller Partner: sh-960-02
Multipath Provider: AIX Native
Multipathing Algorithm: round_robin

-----
MPIO Controller AIX Controller AIX MPIO
path path MPIO host target HBA path
status type path HBA port priority
-----
Enabled primary path0 fcs0 4a 255
Enabled primary path1 fcs0 4b 255
Enabled secondary path2 fcs0 4a 1
Enabled secondary path3 fcs0 4b 1
...
vioserver# sanlun lun show -d hdisk9
    filer: lun-pathname device filename adapter protocol
lun size   lun state
sh-960-01: /vol/vol1/vioserver_7  hdisk9          fcs0      FCP      250m
(262144000)  GOOD
    Serial number: VrHqko8bg72t
    Filer FCP nodename:500a098083e152d9 Filer FCP portname:500a098193e152d9
    Filer adapter name: v.4a
    Filer IP address: 10.60.240.66
    Filer volume name:vol1 FSID:0xe10152d9
    Filer qtree name:/vol/vol1 ID:0x0
```

```
Filer snapshot name: ID:0x0
vioserver#
```

(PowerVM NPIV) Discovering and configuring LUNs

When you are using PowerVM NPIV, all commands are run on the VIO client. This is the same as running AIX MPIO.

Before you begin

(iSCSI) Make sure you have added an entry for the storage system in the file specified by the `disk_filename` parameter.

Steps

1. Log in to the VIO client and enter the following command to discover the newly created LUNs:

```
cfgmgr
```

The `cfgmgr` command causes:

- The initiators to log into the fabric.
- The system to check for new devices.
- The system to create new device entries where the new LUNs are mapped as MPIO devices using the default PCM.

Example

On a VIO server, use the `cfgmgr` command to discover the LUNs.

```
vioclient$ cfgmgr
```

2. To verify that the hdisk MPIO devices were created, enter the following command on the VIO client:

```
lsdev -Cc disk
```

Write down the hdisk instance numbers. You use them when you perform the path configuration in Step 3.

3. Order and prioritize the paths to the storage system LUNs.

- If you are running the FC protocol and a version of Data ONTAP that supports ALUA, you should have ALUA running on the storage system. ALUA automatically sets the paths.

Note: With Data ONTAP 8.0 and later, ALUA is automatically enabled when you create the `igroup`. To confirm that ALUA is enabled, run the command `igroup show -v igrup_name` on the storage system. If it is not enabled, you can use the command `igroup set igrup_name alua yes` to enable it.

- If you are not running ALUA, you must execute the `dotpaths` utility to set the path priorities. The `dotpaths` utility is installed when you install the Host Utilities. If you enter `dotpaths` without any options, it sets the priority for all Data ONTAP LUNs.

Note: Do not use the `dotpaths` utility if ALUA is enabled.

The command line for this utility has the following format:

```
dotpaths [-hqv] [hdiskN ...]
```

You can use the command line options to:

- Specify a set of disks for which you want to set the priority. When you enter a list of hdisk names, the utility sets the priority only for those disks.
- Display all the priorities. Enter the `-v` option to enable verbose mode.
- Query all or individual disk priorities. Enter the `-q` option to perform this task.

Example

The following is an example of using the `dotpaths` utility.

```
vioclient# dotpaths -v
hdisk2 (path 0): Optimized path - Setting priority to 255
hdisk2 (path 1): Optimized path - Setting priority to 255
hdisk2 (path 2): Unoptimized path - Current priority of 1 is correct
hdisk2 (path 3): Unoptimized path - Current priority of 1 is correct
...
Path priority set and/or verified for 22 disks, 86 total paths.
vioserver#
```

4. To verify that the host has discovered the LUNs, enter the following command:

```
sanlun lun show -p
```

Example

This example shows typical output of the `sanlun lun show -p` command for a direct-attached configuration. It also includes the `sanlun lun show -d hdisk9 -v` command to display information about hdisk9. The output from this command has been truncated.

```
# sanlun lun show -p
ONTAP_PATH: sh-960-01:/vol/vol1/vioserver_7
LUN: 7
LUN Size: 250m (262144000)
Host Device: hdisk9
LUN State: GOOD
Controller_CF_State: Cluster Enabled
Controller Partner: sh-960-02
Multipath Provider: AIX Native
Multipathing Algorithm: round_robin

-----
MPIO Controller AIX Controller AIX MPIO
path path MPIO host target HBA path
status type path HBA port priority
-----
Enabled primary path0 fcs0 4a 255
Enabled primary path1 fcs0 4b 255
Enabled secondary path2 fcs0 4a 1
Enabled secondary path3 fcs0 4b 1
...
vioclient# sanlun lun show -d hdisk9
    filer:    lun-pathname      device filename adapter protocol
lun size    lun state
sh-960-01:  /vol/vol1/vioserver_7  hdisk9          fcs0      FCP      250m
(262144000)  GOOD
```

```

Serial number: VrHqko8bg72t
Filer FCP nodename:500a098083e152d9 Filer FCP portname:500a098193e152d9
Filer adapter name: v.4a
Filer IP address: 10.60.240.66
Filer volume name:vol1 FSID:0xe10152d9
Filer qtree name:/vol/vol1 ID:0x0
vioclient#

```

(Veritas DMP) Discovering and configuring LUNs with Veritas

There are several tasks you need to perform when setting up LUNs for a Veritas DMP environment. The following sections provide information on those tasks.

Next topics

- [*\(Veritas DMP\) Discovering LUNs on a host using VxVM and DMP*](#) on page 72
- [*\(Veritas DMP\) Methods for adding a disk to a disk group*](#) on page 73
- [*\(Veritas DMP\) Methods for creating a volume*](#) on page 74
- [*\(Veritas DMP\) Methods for accessing storage on a Veritas volume*](#) on page 75

(Veritas DMP) Discovering LUNs on a host using VxVM and DMP

In environments using VxVM and DMP, you must perform several steps to get the host to discover the LUNs. You do not need to reboot the host to discover the LUNs.

Steps

1. Log in as root on the host.
2. Execute the `cfgmgr` command on the AIX host to discover the newly created LUNs.

The `cfgmgr` command causes

- The initiators to log into the fabric.
- The system to check for new devices.
- The system to create new device entries.

3. Enter `lsdev -Cc disk` to verify that the host has discovered the LUNs.
4. Enter `vxdctl enable` to enable VxVM to recognize each known hdisk.

The `vxdctl` utility recognizes and assigns names to LUNs on the storage system.

5. Verify that the ASL has detected the new devices by entering the `vxdmpadm listenclosure all` command. In the output, verify that the model names of your storage systems are listed as the `ENCLR_NAME` and `ENCLR_TYPE`. Verify that the `ARRAY_TYPE` is either A/A-NETAPP or, if you have ALUA enabled on the igrup, ALUA.

6. Enter the `vxldmpadm listapm all` command to verify that the APM is properly set. If you are using a version of Veritas Storage Foundation prior to 5.1, the APM will be set to active.

- With Veritas Storage Foundation 5.1, the APM will be set to inactive all other will be set to active.
- With versions of Veritas Storage Foundation prior to 5.1, the APM will be set to active.

(Veritas DMP) Methods for adding a disk to a disk group

Adding a disk to a disk group makes the disk space available when you create a VxVM volume. There are several methods you can use to add a disk to a disk group.

First you need to decide how you want to configure the disks. Be sure to have answers to the following questions before you begin:

- Are you adding single or multiple disks?
- Do you want to add disks associated with a specific controller or target?
- Which disks do you want to add?
- Do you want to add the disk to a free disk pool?
- To which disk group will you add the disk?

Disk that you add must be uninitialized or free. If the disk is uninitialized, you must initialize it before you add it to the group.

Below is a summary of the different methods for adding a disk to a disk group. Use the method you prefer. Refer to Veritas documentation for additional information.

Method	Action
VEA	<ol style="list-style-type: none"> Select the disk to add or the disk group to expand. Select Actions and Add Disk to Dynamic Disk Group. Indicate the disk group that you want. Indicate the disks you want to add.
<code>vxdiskadm</code>	<ol style="list-style-type: none"> Enter the <code>vxdiskadm</code> command to activate the <code>vxdiskadm</code> interface. Select option 1 Add or initialize one or more disks and follow screen directions.

Method	Action
CLI	<ul style="list-style-type: none"> • To configure private and public regions on a VxVM disk, enter the <code>vxdisksetup</code> command. • To add one or more disks to a disk group, enter <code>vxdg addisk</code> command.
SMIT	Enter <code>smit vxvmdiskadmadd</code> to add disks using smit

(Veritas DMP) Methods for creating a volume

Veritas volumes span data across multiple disks. There are several methods you can use to create a volume in a Veritas VxVM and DMP environment.

Below is a summary of the different methods for creating a volume. Use the method you prefer. Refer to Veritas documentation for additional information.

Method	Action
VEA	<ol style="list-style-type: none"> 1. Expand the disk group and select the disk group you want. 2. Select Actions and New Volume. 3. Specify the volume characteristics that meet your requirements.
CLI	<ul style="list-style-type: none"> • Use the <code>vxassist</code> command to create a volume: <code>vxassist -g diskgroup make volume_name length attributes</code> • To add one or more disks to a disk group, enter <code>vxdg addisk</code> command.
SMIT	<ol style="list-style-type: none"> 1. Enter the <code>smit vxvm</code> command. 2. Select VxVM Volumes. 3. Select Add a VxVM Volume. 4. Select the layout that meets your requirements.

(Veritas DMP) Methods for accessing storage on a Veritas volume

You can access storage on a Veritas Volume using a raw Veritas volume or a Veritas File System (VxFS). You cannot configure a volume for both raw and VxFS access.

Below is a summary of the different methods for accessing a volume. Use the method you prefer. Refer to Veritas documentation for additional information.

Method	Action
VEA	<ol style="list-style-type: none"> 1. Select the volume that will contain the file system. 2. Specify Actions, File System, and New File System. 3. Specify the file system requirements.
CLI	<ul style="list-style-type: none"> • Use the <code>mkfs</code> command to create a file system: <code>mkfs -V vxfs /dev/vx/rdsk/diskgroup/volume</code> • Use the <code>mount</code> command to mount the file system: <code>mount -V vxfs /dev/vx/dsk/diskgroup/volume mount_point</code>
SMIT	Enter the <code>smit crvxfslvstd</code> command.

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. The following sections provide information on `sanlun`.

Next topics

[Displaying host LUN information with sanlun](#) on page 77

[Explanation of the sanlun lun show -p output](#) on page 78

[Displaying host HBA information with sanlun](#) on page 79

Displaying host LUN information with sanlun

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

Steps

1. Ensure that you are logged in as root on the host.
2. Change to the `/opt/ontap/santools/bin` directory:

```
cd /opt/ontap/santools/bin
```

3. Enter the `sanlun lun show` command to display LUN information. The command has the following format:

```
sanlun lun show [-v] [-d host device filename | all |storagesystem name  
| storagesystem name:storagesystem pathname] -p
```

`-v` produces verbose output.

`-d` is the device option and can be one of the following:

- `host device filename` specifies the special device file name for the disk on the host.
- `all` lists all storage system LUNs attached to the host.
- `storagesystem name` is the host name of the storage system.
- `storagesystem name:storagesystem pathname` lists the path name of the LUN on the storage system.

(Multipathing environments only) `-p` provides information about the primary and secondary paths available to the LUN when you are using multipathing. You cannot use the `-d` option if you use `-p`. Use the following format:

```
sanlun lun show -p [storagesystem name:storagesystem pathname |
storagesystem name | all ]
```

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 the `all` parameter.

For example, you might enter:

- `sanlun lun show -p`
to display a listing of all the paths associated with the LUN. This information is useful if you need to set up path ordering or troubleshoot a problem with path ordering.
- `sanlun lun show -d /dev/hdisk<xx>`
to display the summary listing of the LUN(s) associated with the host device `/dev/hdisk<xx>` where `xx` the number of the hdisk; for example, you might have `hdisk1` or `hdisk4`.
- `sanlun lun show -v all`
to display verbose output for all the LUN(s) currently available on the host.
- `sanlun lun show toaster`
to display a summary listing of all the LUNs available to the host served by the storage system called `toaster`.
- `sanlun lun show toaster:/vol/vol0/lun0`
to display a summary listing of all the LUNs available to the host served by `lun0` on `toaster`.

Note: The `sanlun lun show` command displays only the LUNs that have been discovered by the host. LUNs that have not been discovered by the host are not displayed.

Explanation of the `sanlun lun show -p` output

The `sanlun lun show -p` command provides details for both MPIO stacks and Veritas DMP stacks.

- host path state: Whether the path is enabled or disabled.
- filer path type: Whether the path to the LUN is primary, secondary, or standby.
 - Primary paths communicate directly using the adapter on the local storage system.
 - Secondary paths are proxied to the partner storage system over the cluster interconnect.
 - Standby occurs when the path is being serviced by a partner storage system in takeover mode.
- device filename: The special device file name for the disk on AIX that represents the LUN.
- host HBA: The name of the initiator HBA on the host.
- primary filer port: The port that provides direct access to a LUN. This is always a primary path.
- partner filer port: The port that provides passive path failover. This is always a secondary path. After the failover of a storage system cluster, the `sanlun lun show -p` command reports secondary paths as secondary but enabled, because these are now the active paths.
- **(Veritas DMP) Multipathing policy:** The multipathing policy is one of the following:

- **A/A (Active/Active)**

Veritas DMP uses more than one path concurrently for traffic. The A/A policy indicates that the B ports on the storage system HBAs are in standby mode and become active only during a takeover. This means that there are two active paths to the cluster at any given time.

- **A/A-NETAPP**

Groups of LUNs connected to a single controller failover as a single failover entity. Failover occurs at the controller level and not at the LUN level.

- **ALUA**

A LUN in an ALUA-enabled array can be accessed through both controllers, using optimized and non-optimized paths. The array informs the host of path options, their current state, and state changes. Using this information, the host can determine which paths are optimized. Failover to the non-optimized path occurs only if all the optimized paths fail.

- The primary and secondary controllers are each connected to a separate group of LUNS. If a single LUN in the primary controller's LUN group fails, all LUNs in that group fail over to the secondary controller's passive LUN group. This array type allows load balancing of I/O across multiple primary paths. Refer to the Veritas documentation for details about load balancing.

.

Displaying host HBA information with `sanlun`

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

Steps

1. Ensure that you are logged in as root on the host.
2. Change to the `/opt/ontap/SANToolkit/bin` directory.
3. At the host command line, enter the following command to display host HBA information:

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

`-c` option produces configuration information that you can use to create igroups.

`-v` option produces verbose output.

`all` lists information for all FC adapters.

`adapter name` lists information for a specified adapter.

Example

The following example displays the type of output you see with the `sanlun fcp show adapter` command line. This example simply displays information on the adapters. The next example includes the `-v` option, which provides more information on the adapter.

```
# sanlun fcp show adapter
```

```
fcs0      WWPN:10000000abcde123
fcs1      WWPN:10000000abcde456
```

This example adds the **-v** option to the `sanlun fcp show adapter` command line, which provides more information on the adapter.

```
# sanlun fcp show adapter -v
fcs0      WWPN:10000000abcde123
WWNN:      20000000c936268b
driver name: /usr/lib/drivers/pci/efcdd
model:      df1000f9
model description: FC Adapter
serial number: 1E32909972
hardware version: not available
driver version: 5.1.0.50
firmware version: CD382A1
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 2 GBit/sec
negotiated speed: 1 GBit/sec
OS device name: fcs0

fcs1      WWPN:10000000c9362687
WWNN:      20000000c9362687
driver name: /usr/lib/drivers/pci/efcdd
model:      df1000f9
model description: FC Adapter
serial number: 1E32909970
hardware version: not available
driver version: 5.1.0.50
firmware version: CD382A1
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 2 GBit/sec
negotiated speed: 1 GBit/sec
OS device name: fcs1
#
```

This example uses the **-c** option with the `sanlun fcp show adapter` command line, which provides more information you need to supply when you create an igroup on the storage system.

```
# sanlun fcp show adapter -c

Enter this controller command to create an initiator group for this
system:
igroup create -f -t aix "p570.gdl.tester.ontap.com" 10000000c974921e
10000000c974921f
```

SAN boot LUNs in an AIX MPIO environment

You can set up a SAN boot LUN to work in an AIX MPIO environment that is running the AIX Host Utilities with either the FC or iSCSI protocol.

The method you use for creating a SAN boot LUN and installing a new OS image on it in an AIX MPIO environment can vary, depending on which protocol you are using.

The sections that follow provide steps for configuring a SAN boot LUN and installing a new OS image on it. The examples in these sections use a host running the AIX B 5L. You can apply this procedure to most configurations.

There are also other procedures for creating a bootable LUN, which these sections do not describe. Refer to the appropriate AIX documentation for details about additional configuration methods.

Note:

Solutions and components are qualified on an ongoing basis. To verify that SAN booting is supported in your configuration, see the N series Service and Support Web site at www.ibm.com/systems/storage/network/interophome.html.

Next topics

[*\(AIX MPIO\) Important notes about creating a SAN boot LUN*](#) on page 81

[*\(AIX MPIO, FC\) Creating a SAN boot LUN in an FC environment*](#) on page 82

[*\(AIX MPIO, FC\) Cloning the root volume group of the AIX operating system*](#) on page 86

[*\(AIX MPIO, iSCSI\) Creating a SAN boot LUN in an iSCSI environment*](#) on page 88

[*\(AIX MPIO\) Updating a SAN boot LUN*](#) on page 99

(AIX MPIO) Important notes about creating a SAN boot LUN

Before you begin the process of creating the SAN boot LUN, make sure your system is set up correctly and you have the information you will need to create the boot LUN.

The following list provides information about the prerequisites for creating a boot LUN:

- You are running the correct version and technology level of the operating system and your host and storage system are configured properly.
- You have confirmed that your configuration supports SAN booting by checking the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.
- You have checked the Release Notes to make sure none of the steps have changed since this guide was produced.

- You have determined the size of that boot LUN needs to be. This LUN must be large enough to provide space for the AIX operating system. See the AIX product documentation to determine how much disk space is needed.
- You are logged in as root.
- **(AIX MPIO, FC)** You know the path limitations for the boot LUN. The SAN boot LUN can have a maximum of four paths. The data LUNs do not have this limitation.
- **(AIX MPIO, iSCSI)** You have the iSCSI node name of the host's initiator or HBA. You need this to create an igrup.
- You have the following documentation available to you as you perform the procedure:
 - This guide, which provides a summary of the key steps.
 - The *Data ONTAP Block Access Management Guide*.
 - The AIX product documentation.

(AIX MPIO, FC) Creating a SAN boot LUN in an FC environment

Creating a boot LUN in an FC environment involves several tasks that you must perform on the host and on the storage system. In addition, you must deploy the boot LUN with path restrictions.

Before you begin

Installing directly to a IBM Data ONTAP FC LUN without restricting the number of paths causes one of the following issues:

- A failure occurs at the end of the install.
- The AIX operating system does not boot from the LUN and returns a boot code of 554.

It is a good practice to make sure you have the World Wide Port Names (WWPN) for each of the host's HBAs. Also, you should have already calculated the size of the boot LUN. It must be large enough to provide space for the operating system.

Steps

1. On your storage system, create two igrups. One is for the boot LUN, and the other is for data LUN.

Example

This example shows the igrup for the boot LUN.

```
filerA> igrup show BootLUN
BootLUN (FCP) (ostype: aix):
  10:00:00:00:c9:40:56:23 (logged in on: 5a, 5b, vtic, 4b, 4a)
```

This example shows the igrup for the data LUNs.

```
DataLUNs (FCP) (ostype: aix):
  10:00:00:00:c9:40:56:23 (logged in on: 5a, 5b, vtic, 4b, 4a)
```

```
10:00:00:00:c9:40:56:e9 (logged in on: 5a, 5b, vtic, 4a, 4b)
10:00:00:00:c9:50:54:5a (logged in on: 5a, 5b, vtic, 4a, 4b)
10:00:00:00:c9:50:5f:8c (logged in on: 5a, 5b, vtic, 4a, 4b)
```

2. On the storage system where you created the boot LUN, create a portset for one target port. Ignore the failover port for the boot LUN igroup. Use the following command:

```
portset create -f PortSetName FilerTargetPortName
```

PortSetName is the name of the portset you created; for example, SANboot.

FilerTargetPortName identifies the storage system port, such as controllerA:4a.

Example

This example creates a portset called sanboot for the boot LUN.

```
filerA> portset create -f SANboot controllerA:4a
```

3. Bind the portset to the igroup for the boot LUN using the command

```
igroup bind iGroupName PortSetName
```

iGroupName is the name of the igroup you created; for example, booothost.

PortSetName is the name of the portset you created; for example, sanboot.

Example

This example creates a portset called sanboot associated with the igroup booothost. It then uses the `portset show` command to confirm the port has been set up. Next it uses the `igroup show` command to view the igroup.

```
filerA> igrup bind booothost sanboot
filerA> portset show
    sanboot (FCP):
        ports:
            filerA 0c
        igrups:
            booothost
filerA> igrup show booothost
    booothost (FCP) (ostype: aix):
        10:00:00:00:c9:50:54:5a (logged in on: vtic, 0c)
        10:00:00:00:c9:50:5f:8c (logged in on: 1b, vtic, 0d)
    Pset: sanboot
```

4. Install the AIX operating system on the boot LUN.

Note: For details on installing the AIX operating system, follow the instructions in the appropriate AIX documentation.

5. Install the Host Utilities on the boot LUN according to the instructions in this manual.
6. Discover the devices (IBM Data ONTAP LUNs).
7. Change the boot LUN to avoid SCSI-2 Reserve conflicts (SCSI status 18). The SCSI-2 Reserve is the result of AIX 5L default behavior during the install process.
8. On the host, set the boot LUN to use the AIX MPIO settings.

Note: You do not have to change the settings for data LUNs. Data LUNs can use the Host Utilities default settings.

You want to:

- Identify which device is bound to the boot LUN.
- Change the boot LUN's reserve policy for AIX MPIO to single_path.
- Verify that the `reserve_policy` attribute is set to `single_path` and that the `Algorithm` attribute is set to `fail_over`.

Example

This example uses several commands to set the the AIX MPIO information.

```
AIX_Host> bootinfo -b
hdisk0
AIX_Host> chdev -Pl hdisk0 -a algorithm=fail_over -a
reserve_policy=single_path
AIX_Host> # lsattr -El hdisk0 -a algorithm -a reserve_policy
algorithm      fail_over Algorithm      True
reserve_policy single_path Reserve Policy True
```

9. Restart the AIX server so that it boots from the boot LUN. This enables the MPIO and SCSI-2 Reserve parameters you set in the previous steps.
10. Re-enable the target ports for the boot LUN igroup so that there is a maximum of four paths to the boot LUN.

Note: To avoid a Single Port of Failure (SPOF), avoid disabling both the storage system ports allocated to the portset at the same time or during a takeover event.

Example

The following example re-enables the target ports.

```
filerA> igrup add BootLUN 10:00:00:00:c9:50:5f:8c
filerA> igrup show BootLUN
  BootLUN (FCP) (ostype: aix):
    10:00:00:00:c9:40:56:23 (logged in on: 5a, 5b, vtic, 4b, 4a)
    10:00:00:00:c9:50:5f:8c (logged in on: 5a, 5b, vtic, 4a, 4b)
    Pset: SANboot
filerA> portset add SANboot 4a
filerA> portset show
  SANboot (FCP):
    ports:
      filerA 4a
      filerB 4a
    igrups:
      BootLUN
```

11. On the host, verify that the boot LUN has only the recommended four paths. You can use the `cfgmgr` command to perform this task.

Example

The following example uses the `cfgmgr` command to verify that the boot LUN has an acceptable number of paths. Then it uses the `lspath` command to display the paths.

```
AIX_Host> cfgmgr
AIX_Host> # lspath -l hdisk0
Enabled hdisk0 fscsi0
Enabled hdisk0 fscsi0
Enabled hdisk0 fscsil
Enabled hdisk0 fscsil
```

12. Identify any ghost devices left over from the ODM device overlay from MPIO FC devices or FC devices to Data ONTAP FC devices. You can use the `lsdev` command to perform this task.

Example

The following example shows ghost devices with two different device handles: Defined and Unknown.

```
lsdev -Cc disk
hdisk0 Available 02-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk1 Defined 02-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk2 Defined 03-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk3 Unknown 05-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk4 Unknown 06-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk5 Available 02-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk6 Available 03-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk7 Available 05-08-01 MPIO  ONTAP  FCP Default PCM Disk
hdisk8 Available 06-08-01 MPIO  ONTAP  FCP Default PCM Disk
```

13. Remove the ghost devices.

- For device handles marked Unknown, use the `odmdelete` command.
- For device handles marked Defined, use the `rmdev` command.

Example

Unknown device handles:

```
odmdelete -q name=hdisk3 -o CuDv
0518-307 odmdelete: 1 objects deleted
```

Defined device handles:

```
rmdev -dl hdisk1
hdisk1 deleted
```

14. Rescan the bus for MPIO FC devices. You can use the `cfgmgr` to perform this task.

15. If you are not using ALUA, set the path priorities by running the `dotpaths` utility that comes with the Host Utilities.

Example

Executing the following command line runs the `dotpaths` utility in verbose mode:

```
AIX_Host> /usr/sbin/dotpaths -v
```

16. Invalidate the current bootlist and create a bootlist that boots only from the primary paths of the boot LUN.

Note: This procedure helps prevent a common boot issue where the boot list becomes filled with paths that cannot be used during the boot cycle.

You can use the following commands to perform this procedure:

```
bootinfo -b
bootlist -m normal -i
bootlist -m normal <BootDeviceHandle>
```

These commands perform the following tasks:

- Return information on the last boot device.
- Invalidate the normal mode boot list.
- Set a the new boot device to a normal boot list.

(AIX MPIO, FC) Cloning the root volume group of the AIX operating system

Another way to install the AIX operating system is to clone the root volume. You use the `alt_disk_install` command to clone the volume and then you use `bootlist` and `bootinfo` commands to get information about the new boot device and set it to the boot list.

Steps

1. Enter the `alt_disk_install` command and specify the root volume that you want to clone.

Example

The following command clones hdisk1.

```
-> alt_disk_install -C hdisk1
+-----+
ATTENTION: calling new module /usr/sbin/alt_disk_copy. Please see the
alt_disk_copy man page and documentation for more details.
Executing command: /usr/sbin/alt_disk_copy -d "hdisk1"
+-----+
Calling mkszfile to create new /image.data file.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5.
Creating logical volume alt_hd6.
Creating logical volume alt_hd8.
Creating logical volume alt_hd4.
Creating logical volume alt_hd2.
Creating logical volume alt_hd9var.
Creating logical volume alt_hd3.
Creating logical volume alt_hd1.
Creating logical volume alt_hd10opt.
Creating logical volume alt_sysdump_dev.
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
```

```

Creating /alt_inst/opt file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the alternate file
system...
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
Fixing LV control blocks
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk1

```

2. Reboot the server.
3. Use the `bootlist` command to create a new boot list that boots only from the primary paths of the boot LUN.

Example

The following example contains commands to perform the following tasks:

- Create a boot list in normal mode that says to boot off logical volume hd5 on disk hdisk1 and displays information about the boot.
- Return information on the last boot device.
- Display a 1 if the IPL code in the ROS on the machine running the command is capable of booting from a specified disk.

```

-> bootlist -m normal -o hdisk1 blv=hd5
-> bootinfo -b hdisk1
-> bootinfo -B hdisk1 1

```

4. Use the `ipl_varyon` command to varyon the root volume group.

Example

```

5300-05-00] root@ibmbc-fak01-b4 / -> ipl_varyon -i
PVNAME BOOT DEVICE PVID VOLUME GROUP ID
hdisk0 YES 002006ca49c93e760000000000000000
002006ca0000d000
hdisk1 YES 002006cac6d2a0fe0000000000000000
002006ca0000d000

```

(AIX MPIO, iSCSI) Creating a SAN boot LUN in an iSCSI environment

Creating a boot LUN in an iSCSI environment involves several tasks that you must perform on the host and on the storage system.

About this task

The following provides a high-level look at creating and setting up igroups and LUNs. Detailed information on these tasks is in the *Data ONTAP Block Access Guide*. For more information on the steps that involve installing and setting up your operating system or for other ways to create a SAN boot LUN, see the documentation that came with your AIX system.

Note: To make the examples shown in these steps easier to read, the options you need to select are shown in bold.

Steps

1. Make sure that you are logged in as root and that you have:
 - Calculated the size needed for the boot LUN.
 - Have the iSCSI node name of the host's iSCSI software initiator or HBA.
2. **(Host)** Boot the AIX server using either a NIM server or a CD-ROM. The Base Operating System Installation and Maintenance menu appears.
3. **(Host)** Configure the iSCSI software initiator or iSCSI HBA. Choose option 4 in the Base Operating System Installation and Maintenance menu.

Example

When you start the AIX operating system, it displays the Base Operating System Installation and Maintenance menu. At the “Installation and Maintenance” menu, choose the option: “4 Configure Network Disks (iSCSI).”

```
Welcome to Base Operating System
      Installation and Maintenance
Type the number of your choice and press Enter. Choice is indicated by
>>>.

>>> 1 Start Install Now with Default Settings
      2 Change/Show Installation Settings and Install
      3 Start Maintenance Mode for System Recovery
      4 Configure Network Disks (iSCSI)

88 Help ?
```

```

99 Previous Menu

>>> Choice [1]: 4

Move cursor to desired item and press Enter.

Configure iSCSI -----
Network Utilities

Esc+1=Help Esc+2=Refresh Esc+3=Cancel F8=Image
F9=Shell F10=Exit Enter=Do

x          Available Network Devices      x
x
x Move cursor to desired item and press Enter. x
x
x en0 N/A          x
x ics0 02-09 N/A      x
x

```

4. **(Storage system)** Get information from the storage system on the Data ONTAP controller. You can do this by opening a window on the host and using a command such as telnet to connect to the storage system. You need information on the controller in order to configure the iSCSI adapter for booting.

From a command line on the storage system, you must:

- Get the iSCSI node name. Use the `iscsi nodename` command.
- Display information about the iSCSI portal. Use the `iscsi portal show` command.
- (Recommended)** Disable all the interfaces except the one you are using. Use the `iscsi interface disable interface_name` command.
- Display information on the interface to confirm that the unused interfaces have been disabled. Use the `iscsi interface show` command.
- Create the igrup. Use the `-i -t aix initiator_group` command.
- Display information on the new igrup. Use the command `iscsi show igroup_name` command.

Example

The following is an extended example showing the series of commands that you might execute to get information about the controller. These commands take you through the process of getting the iSCSI node information through setting up an igrup called iSCSI_Boot_POWER6. All these commands are run on the storage system, even though you might be executing them from a window on your host.

```

ControllerA> iscsi nodename
iSCSI target nodename: ign.1986-03.com.ibm:sn.101173314

ControllerA> iscsi portal show
Network portals:
IP address TCP Port TPGroup Interface

```

```

10.60.240.124 3260 1000 e0a

ControllerA> iscsi interface disable e0b e0c e0d

ControllerA> iscsi interface show
Interface e0a enabled
Interface e0b disabled
Interface e0c disabled
Interface e0d disabled

ControllerA> igrup create -i -t aix iSCSI_Boot_POWER6 iqn.1986-
03.com.ibm:sn.101173314

ControllerA> igrup create -i -t aix
  iSCSI_Boot_POWER6 (iSCSI) (ostype: aix):
    iqn.1986-03.com.ibm:sn.106024177 (not logged in)

```

5. (Storage system) Create the boot LUN and map it to the igrup.

From a command line on the storage system, you must:

- Create the boot LUN. You can use the interactive `lun setup` command or the command line `lun create -s size -t aix lun_path`.
- Map the LUN to the igrup you created. Use the `lun map lun_path igrup_name [lun_id]` command.
- Display information on the boot LUN to confirm that it has been created and mapped to the igrup. Use the `lun show igrup_name` command.

Example

The following is an extended example showing the series of commands that you might execute to create a LUN, map it to an igrup, and verify it. All these commands are executed on the storage system, even though you might be executing them from a window on your host using a command such as `telnet`.

```

ControllerA> lun create -s 45g -t aix /vol/SANboot/p570-aix01p1.bootlun

ControllerA> lun map /vol/SANboot/p570-aix01p1.bootlun iSCSI_Boot_POWER6

ControllerA> lun show -v -g iSCSI_Boot_POWER6
/vol/SANboot/p570-aix01p1.bootlun 45g (48318382080) (r/w, online,
mapped)
  Serial#: C4daG4FGMKC/
  Share: none
  Space Reservation: enabled
  Multiprotocol Type: aix
  Maps: iSCSI_Boot_POWER6=0

```

6. (Host) On the host, add the iSCSI controller information to the iSCSI Configuration -- SW Initiator menu panel.

Example

The following example shows the AIX menu screen filled in with the information on the controller that you got during Step 4.

```

iSCSI Configuration -- SW Initiator

Type or select values in entry fields.
Press Enter AFTER making all desired changes.
[Entry Fields]

* Ethernet Interface          en0
* IP Address of Ethernet Interface  10.60.241.77
* Network Mask of Ethernet Interface  255.255.252.0
* Gateway to iSCSI Target      10.60.240.1
* iSCSI Target Name          [iqn.1986-03.com.ibm:s>
* IP Address of iSCSI Target      [10.60.240.124]
* Port Number of iSCSI Target      [3260]
* iSCSI Initiator Name        [.cm.ibm:sn.
106024177]

```

7. **(Host, Storage System)** After you configure the iSCSI initiator, it should log into the storage system.

Example

The following information appears on the storage system when the iSCSI initiator logs into the target.

Note: If the initiator does not login, then there is a problem. Confirm that you got the information on the initiator and entered it correctly.

```

ControllerA> Tue Jan 8 09:56:44 EST [3020-aix07: iscsi.notice:notice]:
iSCSI: New session from initiator
iqn.1986-03.com.ibm:sn.106024177 at IP addr 10.60.241.77

COMMAND STATUS

Command: OK      stdout: yes      stderr: no
Before command completion, additional instructions may appear below.

en0 changed
inet0 changed
iscsii1 changed
hdisk0 Available N/A

```

8. **(Host)** Return to the Base Operating System Installation and Maintenance menu and choose option 2 to change the installation settings and install the operating system.

Example

The following sample screen is the same one you saw when you first started AIX on the host. Select the option: “2 Change/Show Installation Settings and Install.”

```

Welcome to Base Operating System
Installation and Maintenance

Type the number of your choice and press Enter. Choice is indicated by
>>>.

>>> 1 Start Install Now with Default Settings

```

```

2 Change/Show Installation Settings and Install
3 Start Maintenance Mode for System Recovery
4 Configure Network Disks (iSCSI)

88 Help ?
99 Previous Menu

>>> Choice [1]: 2

```

9. **(Host)** Select the iSCSI boot device by choosing option 1 at the System Backup Installation and Settings screen.

Example

The following is an example of the screen you see. Select the option: “1 Disk(s) where you want to install Use Maps.”

```

System Backup Installation and Settings

Either type 0 and press Enter to install with the current settings, or
type
the number of the setting you want to change and press Enter.

Setting:                               Current Choice(s):
1 Disk(s) where you want to install .....      Use Maps..... No
2 Shrink File Systems..... No
3 Import User Volume Groups..... No
4 Recover Devices..... No

>>> 0 Install with the settings listed above.

+-----+
88 Help ? | Selected disks don't match disks with maps.
99 Previous Menu | 

>>> Choice [0]: 1

```

10. **(Host)** Use the iSCSI initiator to choose the device handle.

Note: You can use option 77 to verify that this device handle is via the iSCSI initiator. You must select this option twice.

Example

After you select the option: “1 Disk(s) where you want to install Use Maps,” the following screen appears. Select option 77 twice to verify that the device handle is from the iSCSI software driver.

```

Change Disk(s) Where You Want to Install
Type one or more numbers for the disk(s) to be used for installation
and press
Enter. To cancel a choice, type the corresponding number and Press
Enter.
At least one bootable disk must be selected. The current choice is
indicated
by >>>.

```

```

      Name      Location Code  Size(MB)  VG Status  Bootable  Maps
>>> 1 hdisk0      none          46080    rootvg     Yes      Yes

>>> 0 Continue with choices indicated above
55 More Disk Options
66 Disks not known to Base Operating System Installation
77 Display Alternative Disk Attributes
88 Help ?
99 Previous Menu

>>> Choice [0]: 77 (----- choose option 77 twice to verify this Device
Handle is from the iSCSI Software Driver

      Change Disk(s) Where You Want to Install

Type one or more numbers for the disk(s) to be used for installation
and press
Enter. To cancel a choice, type the corresponding number and Press
Enter.
At least one bootable disk must be selected. The current choice is
indicated
by >>>.

      Name      Device Adapter Connection Location
                  or Physical Location Code
>>> 1 hdisk0  iscsi (----- iSCSI Software Driver

```

11. (Host) Choose option 0 to continue installing to this iSCSI boot device handle.

Example

After you select the option 77, the following screen appears. To continue installing the iSCSI boot device handle, select the option “0 Continue with choices indicated above.”

```

      >>> 0 Continue with choices indicated above
55 More Disk Options
66 Disks not known to Base Operating System Installation
77 Display Alternative Disk Attributes
88 Help ?
99 Previous Menu

>>> Choice [0]: 0

      System Backup Installation and Settings

Either type 0 and press Enter to install with the current settings, or
type
the number of the setting you want to change and press Enter.
      Setting:                                Current Choice(s):

      1 Disk(s) where you want to install ..... hdisk0
          Use Maps..... No
      2 Shrink File Systems..... No
      3 Import User Volume Groups..... No
      4 Recover Devices..... No

>>> 0 Install with the settings listed above.

```

```

88 Help ?          +-----+
99 Previous Menu |   WARNING: Base Operating System Installation will
                   |   destroy or impair recovery of ALL data on the
                   |   destination disk hdisk0.
>>> Choice [0]: 0

```

12. After installing directly to aOntap iSCSI device, the iSCSI boot banner appears followed by the AIX kernel banner. These banners indicate that you have successfully installed the operating system on the boot LUN.

Example

The banners that appear should appear similar to the following ones.

```

iSCSI BOOT -----
Server IP.....10.60.240.124
Client IP.....10.60.241.77
Subnet Mask....255.255.252.0
iSCSI Initiator....iqn.1986-03.com.ibm:sn.106024177
iSCSI Target.....iqn.1986-03.com.ibm:sn.101173314
Target Port.....3260
Target LUN.....0
-----
          Welcome to AIX.
          boot image timestamp: 15:06 01/08
          The current time and date: 15:09:32 01/08/2008
          number of processors: 1 size of memory: 2048MB
boot device:
/lhea@23c00200/ethernet@23e00000:iscsi,ciaddr=10.60.241.77,subnetmask=
255.255.252.0,itname=iqn.1986-
03.com.ibm:sn.106024177,iport=3260,ilun=0,iname=iqn.1986-
03.com.ibm:sn.101173314,siaddr=10.60.240.124,2
          kernel size: 15001975; 64 bit kernel
-----

```

13. (Host) Get AIX to discover the iSCSI boot device. Depending on the controller version being used, AIX discovers the iSCSI boot device as either “MPIO other device” or, if you have an earlier controller, “iSCSI other device.” You use the commands bootlist, lsdev, and ls1pp to discover the boot device.

Example

The following is an example of discovering the iSCSI boot device.

```

[5300-07-01-0748] root@p01p1 / -> bootlist -m normal -o
hdisk0

[5300-07-01-0748] root@p01p1 / -> lsdev | egrep 'ics|iscsi'
ics0          Available 02-09      iSCSI Adapter
iscsi0        Available 02-09-01  iSCSI Protocol Device
iscsil        Available          iSCSI Protocol Device

[5300-07-01-0748] root@p01p1 / -> lsdev -c driver -t iscsi
iscsil Available iSCSI Protocol Device

[5300-07-01-0748] root@p01p1 / -> lsdev -Cc disk | grep -w hdisk0
hdisk0 Available MPIO Other iSCSI Disk Drive

[5300-07-01-0748] root@p01p1 / -> ls1pp -i Ontap*
ls1pp: 0504-132 Fileset Ontap* not installed.

```

```
[5300-07-01-0748] root@p01p1 / -> lsattr -El hdisk0
PCM          PCM/friend/OntapDefaultPCM          Path Control
Module       False
PR_key_value none
Value        True
algorithm   round_robin
Algorithm   True
clr_q       no
error       True
dist_err_pcnt 0
Time        True
dist_tw_width 50
Time        True
hcheck_cmd  inquiry
Command     True
hcheck_interval 30
Interval    True
hcheck_mode  nonactive
Mode        True
host_addr   199.199.199.64
Address     False
location    Location
Label       True
lun_id      0x6000000000000000
ID          False
lun_reset_spt yes
Reset       True
max_transfer 0x80000
Size        True
port_num    0xcbc
Number      False
pvid        00067fbad453c2160000000000000000
identifier  False
q_err       yes
bit         True
q_type     simple
TYPE        True
qfull_dly  2
TASK SET FULL True
queue_depth 64
DEPTH      True
reassign_to 120
value       True
reserve_policy no_reserve
Policy      True
rw_timeout 120
value       True
start_timeout 60
value       True
target_name iqn.1992-08.com.ontap*:sn.151707872 Target
NAME        False
PR_key_value none
Value        True
algorithm   round_robin
Algorithm   True
clr_q       no
error       True
dist_err_pcnt 0
Time        True
dist_tw_width 50
Time        True
hcheck_cmd  inquiry
Command     True
```

hcheck_interval	30	True	Health Check
Interval		True	Health Check
hcheck_mode	nonactive	True	Health Check
Mode		True	Health Check
host_addr	199.199.199.64		Hostname or IP
Address		False	
location			Location
Label		True	
lun_id	0x6000000000000000		Logical Unit Number
ID		False	
lun_reset_spt	yes		LUN Level
Reset		True	
max_transfer	0x80000		Maximum TRANSFER
Size		True	
port_num	0xcbc		PoRT
Number		False	
pvid	00067fbad453c2160000000000000000		Physical volume
identifier		False	
q_err	yes		Use QERR
bit		True	
q_type	simple		Queueing
TYPE		True	
qfull_dly	2		Delay in seconds for SCSI
TASK SET FULL	True		
queue_depth	64		Queue
DEPTH		True	
reassign_to	120		REASSIGN time out
value		True	
reserve_policy	no_reserve		Reserve
Policy		True	
rw_timeout	120		READ/WRITE time out
value		True	
start_timeout	60		START unit time out
value		True	
target_name	iqn.1992-08.com.ontap*:sn.151707872	Target	
NAME		False	

14. (Host) Install the Host Utilities according to the instructions in this manual. For information on doing this

15. (Host) Change the boot LUN's reserve policy for MPIO to single_path.

Note: You do not have to change the settings for data LUNs. Data LUNs can use the Host Utilities default settings.

Example

The following command sets the reserve policy for MPIO to single_path.

```
AIX_Host> chdev -P1 hdisk0 -a algorithm=fail_over -a reserve_policy=single_path
```

16. (Host) Reboot the server.

17. (Host) Verify the boot device is a Ontap* MPIO iSCSI device.

Example

The following is an example of verifying that the boot device is a Ontap* MPIO iSCSI device.

```
[5300-07-01-0748] root@p570-aix01p1 / -> lsllpp -i Ontap*
      Vendor
Fileset   Code      Product Id      Feature Id      Package Name
-----
```

```

Path: /usr/lib/objrepos
Ontap.SAN_toolkit.sanlun
5.1.0.0 Ontap.SAN_toolkit
Ontap.SAN_toolkit.scripts 5.1.0.0
Ontap.SAN_toolkit
Ontap.MPIO_Host_Utils_kit.config 5.1.0.0
Ontap.MPIO_Host_Utils_kit
Ontap.MPIO_Host_Utils_kit.fcp 5.1.0.0
Ontap.MPIO_Host_Utils_kit
Ontap.MPIO_Host_Utils_kit.iscsi 5.1.0.0
Ontap.MPIO_Host_Utils_kit
Ontap.MPIO_Host_Utils_kit.pcmadm 5.1.0.0
Ontap.MPIO_Host_Utils_kit

[5300-07-01-0748] root@p570-aix0lp1 / -> lsdev -Cc disk | grep -w $(bootinfo -b)
hdisk1 Available MPIO Ontap iSCSI Disk
[5300-07-01-0748] root@p570-aix0lp1 / -> lsdev -Cc disk
hdisk0 Defined MPIO Other iSCSI Disk Drive

hdisk1 Available MPIO Ontap. iSCSI Disk

[5300-07-01-0748] root@p570-aix0lp1 / -> rmdev -dl hdisk0
hdisk0 deleted

[5300-07-01-0748] root@p570-aix0lp1 / -> bootlist -m normal -o hdisk1

[5300-07-01-0748] root@p570-aix0lp1 / -> lsattr -El hdisk1
PCM PCM/friend/OntapDefaultPCM Path Control Module
PR_key_value none Persistant Reserve Key
Value True
algorithm round_robin
Algorithm True
clr_q no Device CLEARS its Queue on
error True Distributed Error Sample
dist_err_pcnt 0
Time True Distributed Error Sample
dist_tw_width 50
Time True
hcheck_cmd inquiry Health Check
Command True
hcheck_interval 30 Health Check
Interval True
hcheck_mode nonactive Health Check
Mode True
host_addr 199.199.199.64 Hostname or IP
Address False
location
Label True
lun_id 0x6000000000000000 Logical Unit Number
ID False
lun_reset_spt yes LUN Level
Reset True
max_transfer 0x80000 Maximum TRANSFER
Size True
port_num 0xcbc PORT
Number False
pvid 00067fbad453c21600000000000000000 Physical volume
identifier False
q_err yes Use QERR
bit True
q_type simple Queuing
TYPE True
qfull_dly 2 Delay in seconds for SCSI

```

TASK_SET_FULL	True	
queue_depth	64	Queue
DEPTH		True
reassign_to	120	REASSIGN time out
value	True	
reserve_policy	no_reserve	Reserve
Policy		True
rw_timeout	120	READ/WRITE time out
value	True	
start_timeout	60	START unit time out
value	True	
target_name	iqn.1992-08.com.ontap*:sn.151707872	Target
NAME		False
PR_key_value	none	Persistant Reserve Key
Value	True	
algorithm	round_robin	
Algorithm		True
clr_q	no	Device CLEARS its Queue on
error	True	
dist_err_pcnt	0	Distributed Error Sample
Time	True	
dist_tw_width	50	Distributed Error Sample
Time	True	
hcheck_cmd	inquiry	Health Check
Command		True
hcheck_interval	30	Health Check
Interval		True
hcheck_mode	nonactive	Health Check
Mode		True
host_addr	199.199.199.64	Hostname or IP
Address		False
location		Location
Label		True
lun_id	0x6000000000000000	Logical Unit Number
ID		False
lun_reset_spt	yes	LUN Level
Reset		True
max_transfer	0x80000	Maximum TRANSFER
Size		True
port_num	0xcbc	PORT
Number		False
pvid	00067fbad453c2160000000000000000	Physical volume
identifier		False
q_err	yes	Use QERR
bit		True
q_type	simple	Queuing
TYPE		True
qfull_dly	2	Delay in seconds for SCSI
TASK_SET_FULL	True	
queue_depth	64	Queue
DEPTH		True
reassign_to	120	REASSIGN time out
value	True	
reserve_policy	no_reserve	Reserve
Policy		True
rw_timeout	120	READ/WRITE time out
value	True	
start_timeout	60	START unit time out
value	True	
target_name	iqn.1992-08.com.ontap:sn.151707872	Target
NAME		False

18. (Host) Restart the AIX server so that it boots from the boot LUN with the MPIO and SCSI-2 Reserve parameters you set.

19. (Host) Invalidate the current bootlist and create a bootlist that boots only from the primary paths of the boot LUN.

Note: This procedure helps prevent a common boot issue where the boot list becomes filled with paths that cannot be used during the boot cycle.

You can use the following commands to perform this procedure:

```
bootinfo -b
bootlist -m normal -i
bootlist -m normal <BootDeviceHandle>
```

These commands perform the following tasks:

- Return information on the last boot device.
- Invalidate the normal mode boot list.
- Set a the new boot device to a normal boot list.

(AIX MPIO) Updating a SAN boot LUN

To update a SAN boot LUN to the latest version of the Host Utilities, you must remove the current version of the Host Utilities and then install the new version. You will have to reboot the server after you install the Host Utilities.

About this task

Each version of the Host Utilities is tested to determine the best values for the parameters. As a result, a new version of the Host Utilities may contain updated ODM stanzas to enhance certain parameters. You cannot update the Host Utilities software containing the new ODM stanzas while the boot LUN device handle is open. If you do, the new enhancements will not take effect until the device handle is closed and then reopened. Within a SAN boot environment, the boot LUN device handle cannot be closed until you reboot the AIX server.

Steps

1. Determine the current boot device by executing the `bootinfo -b` command.

Example

This example uses `hdisk161` as the boot LUN. This is a FC LUN from a Data ONTAP storage array.

```
-> bootinfo -b
hdisk161

-> lsdev -Cc disk
hdisk0  Available 01-08-00 SAS Disk Drive
hdisk161 Available 03-02-01 MPIO NetApp FCP Default PCM Disk
```

- Check to see if the Host Utilities for MPIO environments has been installed and, if so, which version is installed. Use the `lslpp` command to get this information.

You cannot remove the boot device handle `hdisk161` because an open device handle cannot be removed. This is why you must reboot the sever after you install the Host Utilities. If you attempt to remove the boot device handle now, the system will display an error message.

```
-> rmdev -dl hdisk161
Method error (/usr/lib/methods/ucfgdevice):
      0514-062 Cannot perform the requested function because the
      specified device is busy.
```

- Remove the Host Utilities software using either the command `installpp -u FileSetName` SMIT.
- Confirm that the boot device handle has been orphaned by executing the `lsdev -Cc disk` command.

Example

This example shows that the `hdisk161` has been orphaned and the customized configuration is no longer available.

```
-> lsdev -Cc disk
lsdev: 0514-521 Cannot find information in the predefined device
      configuration database for the customized device hdisk161.
hdisk0  Available 01-08-00 SAS Disk Drive
hdisk161 Available 03-02-01 N/A
```

- Install the Host Utilities software using either the `installpp` command or SMIT.

Example

```
-> installpp -ld .
Fileset Name          Level          I/U Q Content
=====
ontap.MPIO_Host_Utilities_kit.config 5.1.0.0      I  b  usr
#   Network Appliance MPIO PCM Kit Path Config Tool

ontap.MPIO_Host_Utilities_kit.fcp  5.1.0.0      I  b  usr
#   Network Appliance MPIO PCM Kit FCP Disk ODM Stanzas

ontap.MPIO_Host_Utilities_kit.iscsi 5.1.0.0      I  b  usr
#   Network Appliance MPIO PCM Kit iSCSI Disk ODM Stanza

ontap.MPIO_Host_Utilities_kit.pcmodm 5.1.0.0      I  b  usr
#   Network Appliance MPIO PCM Kit PCM Friend ODM Stanza
```

```
[5300-11-01-0944] root@ibmbc-aix02b6 /tmp/ontap_aix_host_utilities_5.1/MPIO ->
installpp -aXYd . Ontap.MPIO_Host_Utilities_kit
+-----+
          Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...
SUCCESSES
-----
```

```

Filesets listed in this section passed pre-installation verification
and will be installed.

Selected Filesets
-----
ontapp.MPIO_Host_Utils_kit.config 5.1.0.0      # Network Appliance MPIO
PCM K...
ontap.MPIO_Host_Utils_kit.fcp 5.1.0.0          # Network Appliance MPIO
PCM K...
ontap.MPIO_Host_Utils_kit.iscsi 5.1.0.0          # Network Appliance MPIO
PCM K...
ontap.MPIO_Host_Utils_kit.pcmodm 5.1.0.0          # Network Appliance MPIO
PCM K...

<< End of Success Section >>

:
:
:
Finished processing all filesets. (Total time: 2 secs).

installlp: bosboot verification starting...
installlp: bosboot verification completed.
installlp: bosboot process starting...

bosboot: Boot image is 41489 512 byte blocks.
0503-292 This update will not fully take effect until after a
          system reboot.

* * * A T T E N T I O N * * *
System boot image has been updated. You should reboot the
system as soon as possible to properly integrate the changes
and to avoid disruption of current functionality.

installlp: bosboot process completed.
+-----+
                         Summaries:
+-----+
Installation Summary
-----
Name          Level      Part      Event      Result
-----
ontap.MPIO_Host_Utils_kit.pcmo 5.1.0.0      USR      APPLY
SUCCESS
ontap.MPIO_Host_Utils_kit.iscs 5.1.0.0      USR      APPLY
SUCCESS
ontap.MPIO_Host_Utils_kit.fcp  5.1.0.0      USR      APPLY
SUCCESS
ontap.MPIO_Host_Utils_kit.conf 5.1.0.0      USR      APPLY
SUCCESS

```

6. Verify that the boot device is a Data ONTAP FC AIX MPIO capable device by executing the `lsdev -Cc disk` command.

Example

This example confirms that the boot device hdisk161 is correct.

```

-> lsdev -Cc disk
hdisk0  Available 01-08-00 SAS Disk Drive
hdisk161 Available 03-02-01 MPIO NetApp FCP Default PCM Disk

```

7. Set the reserve policy to single_path.

Example

This example sets the reserve policy to single_path.

```
-> chdev -Pl hdisk161 -a "algorithm=fail_over reserve_policy=single_path"
hdisk161 changed
[5300-11-01-0944] root@ibmbc-aix02b6 / -> lsattr -El hdisk161 -a "algorithm
reserve_policy"
algorithm      fail_over      Algorithm      True
reserve_policy  single_path  Reserve Policy  True
```

8. It is a good practice to verify the boot path list before you reboot the server. You can use the bootlist command to check the path list.

Example

This is an example of checking the boot path list.

```
-> bootlist -m normal -o
hdisk161 blv=hd5
hdisk161 blv=hd5
hdisk161 blv=hd5
hdisk161 blv=hd5
hdisk161 blv=hd5
```

9. Reboot the server so that the boot device handle is closed and then re-opened at boot time.

Features supported by the Host Utilities

The Host Utilities support a number of features available with AIX hosts and storage systems running Data ONTAP. Your specific environment affects what the Host Utilities support.

Some of the supported features include:

- **(AIX MPIO)** FC and iSCSI within the same AIX host and to the same target as long as they use different volume groups
- Multiple paths to the storage system when a multipathing solution is installed
- **(AIX MPIO)** Switch trunking to enable the use of iSCSI in environments that do not use AIX's native MPIO for multipathing
- Dynamic Logical Partitioning
- **(FC)** ALUA
- **(FC)** Cluster failover using single image mode only
- PowerVM Live Partition Mobility on IBM System p
- IBM General Parallel File System (GPFS)
- SAN booting
- **(PowerVM)** VIO server/client environment
- **(Veritas DMP)** Veritas Storage Foundation

For information on which features are supported with which environments, see the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.

Next topics

[*\(AIX MPIO\) Multiple protocols supported on a single host*](#) on page 104

[*HBAs and the AIX Host Utilities*](#) on page 104

[*Multipathing and the AIX Host Utilities*](#) on page 104

[*Dynamic LPAR is available*](#) on page 104

[*GPFS support*](#) on page 105

[*Live Partition Mobility from PowerVM 1.5.2*](#) on page 105

[*\(FC\) ALUA support with certain versions of Data ONTAP*](#) on page 105

[*SAN booting and the Host Utilities*](#) on page 107

[*Support for non-English versions of AIX operating systems*](#) on page 107

(AIX MPIO) Multiple protocols supported on a single host

The Host Utilities allow you to have both FC and iSCSI pointing to the same target on the same host as long as you do not have both protocols on the same volume group.

To use both protocols in this way, each protocol must have its own volume group. This means you must use a separate LUN for each protocol. You cannot mix FC and iSCSI on the same volume group.

HBAs and the AIX Host Utilities

The Host Utilities support a number of HBAs.

For details on which HBAs are supported, see the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.

The supported HBAs should be installed before you install the Host Utilities. See your HBA documentation for information on installing the HBAs.

Multipathing and the AIX Host Utilities

The AIX Host Utilities support different multipathing solutions based on your configuration.

Having multipathing enabled allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues on the remaining paths.

The Veritas environment of the Host Utilities uses Veritas DMP to provide multipathing.

The MPIO and PowerVM environments of the Host Utilities use AIX's native multipathing solution (MPIO).

Note: In iSCSI environments, you can use switch trunking to enable the use of iSCSI if your environment does not use MPIO for multipathing.

You can use the Host Utilities `sanlun` command to display the path policy that the host has access to.

Dynamic LPAR is available

The Host Utilities support Dynamic Logical Partitioning (DLPAR). DLPAR is an AIX 5.2 and 5.3 feature on IBM pSeries servers.

With DLPAR, you can move hardware resources such as processors, memory, and I/O slots from one logical partition to another without disruption. Refer to the AIX documentation for details about DLPAR.

GPFS support

Certain Host Utilities configurations support GPFS.

GPFS is an AIX feature that allows parallel applications simultaneous access to the same files, or different files, from any node that has the GPFS file system mounted. At the same time GPFS maintains a high level of control over all file system operations.

For information on which configurations support this feature, see the IBM NAS support page at www.ibm.com/storage/support/nas/.

For information on setting up GPFS, check your operating system documentation.

Note: At the time this document was prepared, IBM provided a GPFS library of documents at <http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.gpfs.doc/gpfsbooks.html>.

Live Partition Mobility from PowerVM 1.5.2

Certain Host Utilities environments support PowerVM Live Partition Mobility on IBM System pSeries servers.

PowerVM Live Partition Mobility on IBM System pSeries servers is an AIX feature that lets you quickly migrate currently running AIX partitions and their hosted applications from one physical server to another without disrupting the infrastructure services. The migration operation maintains complete system transactional integrity as it transfers the system environment, including processor state, memory, attached virtual devices, and connected users.

For information on which configurations support Live Partition Mobility, see the IBM NAS support page at www.ibm.com/storage/support/nas/.

For information on setting up Live Partition Mobility, check your operating system documentation.

Note: At the time this document was prepared, IBM provided a PowerVM Live Partition Mobility library of documents at <http://www14.software.ibm.com/webapp/set2/sas/f/pm/component.html>.

(FC) ALUA support with certain versions of Data ONTAP

The AIX Host Utilities running with the FC protocol support both storage systems that have ALUA enabled for active/active storage controllers (clustered storage systems) and those that do not use ALUA. To use ALUA, you must be running a version of Data ONTAP that supports it.

ALUA defines a standard set of SCSI commands for discovering path priorities to LUNs on FC SANs and managing those paths. When you have the host and storage controller configured to use ALUA, it automatically determines which target ports provide optimized (primary) and unoptimized (secondary) access to LUNs.

Note: ALUA is not supported with iSCSI.

Stand-alone storage controllers provide symmetric access to LUNs and do not use ALUA.

To determine if ALUA is available for your environment, check your version of Data ONTAP to see if it supports ALUA and check the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html to see if the Host Utilities support that version of Data ONTAP.

(FC) Specific AIX Host Utilities environments that support ALUA

You can use ALUA if you have an AIX Host Utilities FC environment running either AIX MPIO or Veritas Storage Foundation 5.1 and a version of Data ONTAP that supports ALUA.

The following AIX environments support ALUA when you are using the FC protocol and Data ONTAP 7.3.1 or later:

Host Utilities version	Host requirements
AIX Host Utilities 4.0, 4.1, 5.0, or 5.1 running on a system using either MPIO or Veritas Storage Foundation 5.1	<p>One of the following or later:</p> <ul style="list-style-type: none"> • 5.2 TL8 • 5.3 TL9 SP4 with APAR IZ53157 • 5.3 TL10 SP1 with APAR IZ53158 • 6.1 TL2 SP4 with APAR IZ53159 • 6.1 TL3 SP1 with APAR IZ53160 <p>Note: It is strongly recommended that, if you want to use ALUA, you use the latest levels of 5.3 TL9 or 6.1 TL2 listed in the support matrix. ALUA is supported on all AIX Service Streams that have the corresponding APAR (authorized program analysis report) installed. At the time this document was prepared, the Host Utilities supported AIX Service Streams with the APARs listed above as well as with APARs IZ53718, IZ53730, IZ53856, IZ54130, IZ57806, and IZ61549. If an APAR listed here has not been publicly released, contact IBM and request a copy.</p>

Note: ALUA is not supported with the iSCSI protocol.

SAN booting and the Host Utilities

The Host Utilities support SAN booting in FC, iSCSI, and Veritas Storage Foundation 5.1 environments. SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an AIX 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.

If you are using the Host Utilities with Veritas Storage Foundation 5.1, follow the instructions provided by Symantec to set up SAN booting. You do not need to perform any storage-specific steps. At the time this guide was prepared, Symantec provided the document *Veritas Volume Manager: Configuring DMP on AIX for SAN Booting*, which was located at ftp://exftp.symantec.com/pub/support/products/Foundation_Suite/308527.pdf when this document was prepared.

Note: While SAN booting is not supported with Jumbo frames at this time, solutions and components are qualified on an ongoing basis. To determine whether SAN booting is supported in your configuration, see the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html.

Support for non-English versions of AIX operating systems

AIX Host Utilities are supported on all language versions of AIX. All product interfaces and messages are displayed in English; however, all options accept Unicode characters as input.

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.

Next topics

[Check the Release Notes](#) on page 109

[If you cannot resolve your problem](#) on page 109

[About the troubleshooting sections that follow](#) on page 110

[FC troubleshooting tips](#) on page 116

[iSCSI troubleshooting tips](#) on page 117

[Basic operations take a long time to complete](#) on page 119

[Installing the nSANity data collection program](#) on page 120

Check the Release Notes

The *Release Notes* contain the most up-to-date information on known problems and limitations. The *Release Notes* also contain information on how to look up information about known issues.

The *Release Notes* are updated when there is new information about the Host Utilities. It is a good practice to check the *Release Notes* before you install the Host Utilities and any time you encounter a problem with the Host Utilities.

If you cannot resolve your problem

If you encounter a problem that you cannot resolve and you do not find a solution in this document or the *Release Notes*, you should contact IBM Support at 1-800- IBM-SERV for assistance.

Make sure you have recorded any error messages displayed on the host or storage system console and the symptoms of the problem and any changes made to the host or storage system just before the problem appeared. Also collect the host and storage system log files.

About the troubleshooting sections that follow

The troubleshooting sections that follow help you verify your system setup. Having the correct system setup resolves the most problems. These sections also provide tips for diagnosing and fixing issues.

If you have any problems with the Host Utilities, make sure your system setup is correct.. As you go through the following sections, keep in mind:

- For more information on AIX commands, see the man pages and operating system documentation.
- For more information on the Data ONTAP commands, see the Data ONTAP documentation, in particular, the *Data ONTAP Block Access Management Guide for iSCSI and FC* for your version of Data ONTAP.
- You perform some of these checks from the host and others from the storage system. In some cases, you must have the Host Utilities SAN Toolkit installed before you perform the check. For example, the SAN Toolkit contains the `sanlun` command, which is useful when checking your system.
- To make sure you have the current version of the system components, see the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html. Support for new components is added on an on-going basis. This online document contains a complete list of supported HBAs, platforms, applications, and drivers.

Next topics

[Possible connectivity problems](#) on page 110

[Storage system checklist](#) on page 111

[Host checklist](#) on page 111

[Useful commands and utilities for troubleshooting](#) on page 112

[Check the version of your host operating system](#) on page 113

[Check the version of the prerequisite filesets](#) on page 113

[Confirm that the HBA is supported](#) on page 114

[Check the LUN characteristics](#) on page 114

[Check the MPIO disk path information](#) on page 115

[Reinstalling if the installation fails due to system requirements](#) on page 116

Possible connectivity problems

If your system encounters an I/O error, read/write failures can occur that causes applications to hang, crash, or fail to return requested information.

You can minimize I/O application error on AIX systems by using highly available, multipathing configurations that provide continuous I/O with no single point of failure. If you continue to experience I/O errors, check the following areas for connectivity problems:

- Switch setup (if used)
- Storage system setup
- Host setup
- Storage system connectivity to the switch
- Storage system connectivity to the host
- FC or iSCSI connectivity to the host

Storage system checklist

There are several things you can check that pertain to your storage system to ensure that your Host Utilities environment works correctly.

Make sure that the following situations are true:

- The host and the initiator HBAs or the software initiators are set up.
- The storage system and target HBAs are configured.
- The storage system is set up to work with the host:
 - **(FC)** The storage has been cabled to the fabric according to your configuration topology.
 - **(iSCSI)** The iSCSI HBA or software initiator is logged onto the storage system.
- You have created working volumes and qtrees (if desired) on the storage system.
- Your protocol (either FC or iSCSI) is licensed and running.
You can check the status of the FC or iSCSI protocol service with the `fcp status` command or the `iscsi status` command.
- All the target ports are in a good state.
- The igroups have the ostype set to AIX.
- The LUNs are
 - mapped to host igroups.
 - in a good state as indicated by the output from the `samlun lun show -p` command or the `lspath` command.
 - using AIX as their operating system type.

Host checklist

If you want to ensure that your Host Utilities environment is operating correctly, there are several things you can check.

Make sure that the following situations are true:

- The system meets the Host Utilities requirements.
The installation fails if your system does not meet the requirements.
- The host and the storage system can communicate. One way to check this is to log in to the storage system remotely from the host. Then enter any command at the storage system prompt to confirm that the host and the storage system can communicate.
- **(FC)** The HBAs are working.

Check the output from the `sanlun fcp show adapter -v` command.

- **(iSCSI)** The iSCSI HBA is logged onto the storage system.

Check the output of either the `igroup show` command or the `iscsi initiator show` command.

Note: There must be open volume groups for these commands to display iSCSI HBA information.

- The host is seeing all of the LUNs and LUN paths that it should.

You can confirm this by running the `lsdev` command or, if you have multipathing, `sanlun lun show -p`.

- **(MPIO)** The paths priorities have been set. If your environment does not use ALUA, you must run the `dotpaths` utility to set the path priorities after having the host discover the LUNs.

- **(MPIO)** Use MPIO commands to check your path state and device availability.

- **(Veritas DMP)** Use Veritas commands to check your path state and device availability. You can also use ALUA to set the path priorities.

Useful commands and utilities for troubleshooting

There are several commands and utilities you can use to check your system and look for problems.

The following commands and utilities are especially useful in tracking down problems:

- The Host Utilities `sanlun` utility
- The AIX `lsdev` command
- **(Veritas environments)** The Veritas commands `vxdmpadm` and `vxdisk list`

(sanlun utility) With the `sanlun` utility, you can perform the following tasks:

- Check the status of the
 - HBA
 - LUN
 - Paths to the LUN
- Discover whether a path is up or down

When you have a direct-attached configuration, use the `sanlun lun show` command to get information about your system.

With a multipathing configuration, use the `sanlun lun show -p` command. This version of the command also provides information about the paths to the LUNs.

(lsdev -Cc disk command) The `lsdev -Cc disk` command provides information about the devices and LUNs in your configuration. The output from this command differs slightly depending on whether you run it in a Veritas environment or an MPIO environment.

Check the version of your host operating system

Make sure you have the correct version of the operating system.

You can use the `oslevel -s` command to display information about your operating system.

This command displays basics information your operating system level and maintenance or technology level. The information may not be complete; you may need additional filesets.

The following command shows the information on your operating system level.

```
#-> oslevel -s  
7100-00-02-1041
```

Check the version of the prerequisite filesets

Make sure you have the correct version of the filesets for your operating system installed.

You can use the `lslpp -L` command to display information about all the filesets on the host.

To get information about a specific fileset, use the `lslpp -L fileset_name` command, where *fileset_name* is the name of the fileset you want to check.

The following command line displays information about the fileset `devices.pci.df1000f7.com`.

```
# lslpp -L devices.pci.df1000f7.com  
Fileset Level State Type Description  
(Uninstaller)  
-----  
devices.pci.df1000f7.com 5.2.0.40 C F Common PCI FC Adapter  
Device Software  
State codes:  
A -- Applied.  
B -- Broken.  
C -- Committed.  
E -- EFIX Locked.  
O -- Obsolete. (partially migrated to newer version)  
? -- Inconsistent State...Run lppchk -v.  
Type codes:  
F -- Installpp Fileset  
P -- Product  
C -- Component
```

T	-- Feature
R	-- RPM Package

Confirm that the HBA is supported

The supported HBAs should be installed before you install the Host Utilities software package.

You can use the `lscfg -vpl adapter_name` command to display information about the HBAs. Then you can compare that information with the information in the NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html to determine whether your HBAs are supported.

For information about the type adapter installed, you can use the `lsdev -Cc adapter` command.

Note: If you have an HBA that is not supported, you will need to install a supported HBA. See the documentation that accompanies the HBA for installation instructions.

In this example, the `lsdev -Cc adapter` command displays information on the HBAs. Both `fcs0` and `fcs1` are Fibre Channel adapters.

```
# lsdev -Cc adapter
ent0          Available 03-08 2-Port 10/100/1000 Base-TX PCI-X
Adapter (14108902)
ent1          Available 03-09 2-Port 10/100/1000 Base-TX PCI-X
Adapter (14108902)
fcs0          Available 05-08 FC Adapter
fcs1          Available 09-08 FC Adapter
ide0          Available 07-08 ATA/IDE Controller Device
sisscsia0    Defined 04-08 PCI-X Dual Channel Ultra320 SCSI Adapter
sisscsial    Available 08-08 PCI-X Dual Channel Ultra320 SCSI Adapter
usbhco       Available 02-08 USB Host Controller (33103500)
usbhc1       Available 02-09 USB Host Controller (33103500)
vsa0          Available      LPAR Virtual Serial Adapter
vsal          Available      LPAR Virtual Serial Adapter
```

Check the LUN characteristics

You can display the characteristics of a LUN by executing the `lsattr -El <hdiskname>` command.

This example uses the `lsattr -El` to display information about `hdisk4`.

```
# lsattr -El hdisk4
PCM          PCM/friend/Default          PCM Path Control
Module       False
algorithm    round_robin
clr_q        no
Queue on error True
dist_err_pcpt 0
Sample Time  True
Algorithm True
Device CLEARS its
Distributed Error
```

dist_tw_width	50	Distributed Error
Sample Time	True	
hcheck_cmd	inquiry	Health Check
Command	True	
hcheck_interval	30	Health Check
Interval	True	
hcheck_mode	nonactive	Health Check
Mode	True	
host_addr	10.60.244.35	Hostname or IP
Address	False	
location		Location
Label	True	
lun_id	0x0	Logical Unit
Number ID	False	
max_transfer	0x40000	Maximum TRANSFER
Size	True	
port_num	0xcbc	PORT
Number	False	
pvid	000a105418b0c7ed0000000000000000	Physical volume
identifier	False	
q_err	yes	Use QERR
bit	True	
q_type	simple	Queuing
TYPE	True	
queue_depth	12	Queue
DEPTH	True	
reassign_to	120	REASSIGN time out
value	True	
reserve_policy	no_reserve	Reserve
Policy	True	
rw_timeout	120	READ/WRITE time
out value	True	
start_timeout	60	START unit time
out value	True	
target_name	iqn.1992-08.com.ontap:sn.33610045	Target
NAME	False	

Check the MPIO disk path information

You can display information about the MPIO disk paths by executing the `lspath` command.

This example uses the `lspath` to display information about the MPIO disk path for an iSCSI-attached device.

```
# lspath -l hdisk20 -F connection:status:parent:path_id
iqn.1986-03.com.ibm:sn.84213404,192.168.2.221,0xcbc,
0x1000000000000:Enabled:iscsi0:0
```

```
ign.1986-03.com.ibm:sn.84213404,192.168.2.221,0xcbc,
0x1000000000000:Enabled:iscsi1:1
```

Reinstalling if the installation fails due to system requirements

The installation process requires certain minimum system requirements. If the process finds a problem with your system requirements, the installation fails and you must reinstall the Host Utilities.

About this task

Complete the following steps to reinstall the Host Utilities.

Steps

1. Ensure that you are logged in as root.
2. Check the installation logs to determine why the install failed.

The most common problems include missing the required filesets, such as devices.common.IBM.fc.hba-api fileset.

3. Correct the problem described in the log.
4. Reinstall the software. You can use SMIT or the `installp` command.

Since the installation failed, you do not need to uninstall the software.

FC troubleshooting tips

The following sections provide tips for working in an MPIO environment that is running the FC protocol.

(FC) Check the ALUA settings on the storage system

If you are using an MPIO environment with FC and you are running a version of Data ONTAP that supports ALUA, you should have ALUA set on the storage system to work with igroups.

In Veritas DMP environments, you can use either VxVM or ALUA.

To verify that you have ALUA enabled for your igroup, execute the `igroup show -v` command.

Note: If you do not have ALUA set up, but you have already installed the Host Utilities, you must remove the existing LUN handles for the device before you can enable ALUA. Use the `rmdev -dl device_name`. For example, if you needed to remove the hdisk2 LUN handles, you would enter the command

```
rmdev -dl hdisk2
```

```
.
```

In this example, executing `igroup show -v` shows that ALUA is enabled.

```
filerA# igrup show -v
filerA (FCP):
OS Type: aix
Member: 10:00:00:00:c9:4b:e3:42 (logged in on: 0c)
Member: 10:00:00:00:c9:4b:e3:43 (logged in on: vtic)
ALUA: Yes
```

iSCSI troubleshooting tips

The following sections provide tips for working in an MPIO environment that is running the iSCSI protocol.

Next topics

[\(MPIO, iSCSI\) Verify information about the HBA](#) on page 117

[\(MPIO, iSCSI\) Verify that the HBA is logged into the storage system](#) on page 118

[\(MPIO, iSCSI\) Verify the adapter's configuration](#) on page 118

(MPIO, iSCSI) Verify information about the HBA

If you are using an MPIO environment with iSCSI, you may want to verify that the HBA is supported.

You can use the `lscfg -vl ics*` command verify the feature code, part number, and firmware level of the IBM iSCSI HBA.

The HBA stays logged into the iSCSI target as long as the iSCSI device handles are currently in an active volume group. Otherwise, the iSCSI device handles are closed, and the HBA logs out of the target after device discovery.

If the HBA is installed in an unsupported server, you may receive an error message when you run the `cfgmgr` command.

In this example, the `lscfg -vl ics*` command displays information about the HBA.

```
# lscfg -vl ics*
  ics0 U787B.001.DNW872B-P1-C1-T2 iSCSI Adapter
  IBM iSCSI Server TX Adapter :
  Part Number.....30R5219
  EC Level.....J85591
  Manufacture ID.....52P8480
  FRU Number.....30R5209
  Customer Card ID Number.....1986 (feature code)
  Brand.....H0
  Serial Number.....YJ1H7N62F08W
  Network Address.....000D60B96D55
```

```
ROM Level.(alterable).....000002000057
Product Specific.(Z1).....iqn.1996-04.com.ibm:sn.0310b806g1
Device Specific.(YL).....U787B.001.DNW872B-P1-C1-T2
```

If the HBA had been installed in an unsupported server, you would see the following message when you ran the `cfgmgr` command.

```
# cfgmgr -l ics0
Method error (/usr/lib/methods/cfgqlisc -l ics0 ): 0514-086
Cannot perform the requested function because the specified
device is not supported on this platform
```

(MPIO, iSCSI) Verify that the HBA is logged into the storage system

If you are using an MPIO environment with iSCSI, you may want to go to the storage system and verify that the HBA is logged onto it.

Note: The HBAs show up as logged in only when there are open volume groups.

You can run either the `igroup show` command or the `iscsi initiator show` command from the storage system to verify that the HBA is logged in.

You can also use the `ping` command to determine if the storage system can communicate with the HBA.

Note: If you are using an HBA, do not use the `ping` command on the host. The HBA does not have an Ethernet interface.

In this example, the `igroup show` command and the `iscsi initiator show` command both indicate that the HBA is logged onto the storage system.

```
*> igrup show
ibmp550-fak01 (iSCSI) (ostype: aix):
iqn.1996-04.com.ibm:sn.0310b806g1 (logged in on: e0)
iqn.1996-04.com.ibm:sn.0310b806g2 (logged in on: e0)
*> iscsi initiator show
Initiators connected:
TSIH TPGGroup Initiator
 35 1000 iqn.1996-04.com.ibm:sn.0310b806g1 /
40:00:0f:21:00:00
 36 1000 iqn.1996-04.com.ibm:sn.0310b806g2 /
40:00:0f:21:00:00
```

(MPIO, iSCSI) Verify the adapter's configuration

If you are using an MPIO environment with iSCSI, you may want to verify that the software initiator or HBA characteristics were configured correctly.

You can use the `lsattr -El icsX` command where `X` is the adapter instance to verify the configuration.

In this example, the `lsattr -El ics` command verifies the adapter instance 0.

```
# lsattr -El ics0
bus_intr_lvl    278                      Bus interrupt
level           False
bus_io_addr     0xcfc00                   Bus I/O
address          False
bus_mem_addr    0xc0100000                Bus memory
address          False
disc_filename   /etc/iscsi/targetshw     Discovery
Filename         True
disc_policy     file                     Discovery
Policy          True
gateway_addr    Adapter                  Gateway
Address          True
host_addr       10.60.224.136           Adapter IP
Address          True
initiator_name  iqn.1996-04.com.ibm:sn.0310b806g1 iSCSI Initiator
Name            True
intr_priority   3                        Interrupt
Priority         False
lg_term_dma     0x800000 Amount of memory used for long term DMA
mapping          True
max_xfer_size   0x100000                 Maximum transfer
size            True
num_cmd_elems   200 Maximum number of commands to queue to the
adapter         True
subnet_mask     255.255.255.0          Adapter Subnet
Mask            True
```

Basic operations take a long time to complete

Sometimes basic operations take hours instead of minutes to complete even though there are no errors and everything appears to be correctly connected

This could mean that one of the connectors is not firmly attached. This can be the case even when the connector has a link and is reporting information, such as the correct speed.

If you encounter this problem, try disconnecting the connectors and reconnecting them. Check for faulty optical cabling. If you are using FC, check the switch for excessive line errors.

Installing the nSANity data collection program

Download and install the nSANity Diagnostic and Configuration Data Collector program when instructed to do so by your technical support representative.

Before you begin

The nSANity program replaces the diagnostic programs included in previous versions of the Host Utilities. The nSANity program runs on a Windows or Linux system with network connectivity to the component from which you want to collect data. It does not run directly on the AIX host itself.

About this task

You can install the product software from software updates available for download. Downloads are available only to entitled IBM N series customers who have completed the registration process on the IBM NAS support Web site.

Steps

1. Check the publication matrix page for important alerts, news, interoperability details, and other information about the product before beginning the installation.
2. Obtain the product software by downloading the software as follows:
 - a. Go to the IBM NAS support Web site.
 - b. Sign in with your IBM ID and password.If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration, click the **Register** link, follow the online instructions, and then sign in.
- c. Select the N series software you want to download, and then select the **Download** view. Select the Windows zip or Linux tgz version of the nSANity program, depending on the workstation or server you want to run it on.
- d. Use the **Software Packages** link on the Web page presented and follow the online instructions to download the software.
3. Change to the directory to which you downloaded the zip or tgz file.
4. Extract all of the files and follow the instructions in the README.txt file. Also be sure to review the RELEASE_NOTES.txt file for any warnings and notices.

After you finish

Run the specific nSANity commands specified by your technical support representative.

Related information

IBM NAS support web site - www.ibm.com/storage/support/nas

Where to find more information

For additional information about host and storage system requirements, supported configurations, best practices, your operating system, and troubleshooting, see the documents listed in the following table.

If you need more information about...	Go to...
Known issues, troubleshooting, operational considerations, and post-release developments	<p>The latest Host Utilities <i>Release Notes</i></p> <p>Note: The <i>Release Notes</i> are updated more frequently than the rest of the documentation. You should always check the <i>Release Notes</i> before installing the Host Utilities to see if there have been any changes to the installation or setup process since this document was prepared. You should check them periodically to see if there is new information on using the Host Utilities. The <i>Release Notes</i> provide a summary of what has been updated and when.</p>
The latest supported configurations	<ul style="list-style-type: none">• The Interoperability Matrix• System Configuration Guide
A summary of some of the commands you might use with the Host Utilities	The <i>Host Utilities Quick Command Reference</i> for your protocol
Changes to the host settings that are recommended by the Host Utilities	<i>Host Settings Affected by the Host Utilities</i>
Configuring the storage system and managing SAN storage on it	<ul style="list-style-type: none">• Data ONTAP documentation Index• <i>Best Practices for Reliability: New System Installation</i>• <i>Data ONTAP Software Setup Guide</i>• <i>Data ONTAP Block Access Management Guide for iSCSI and FC</i>• <i>Data ONTAP Release Notes</i>• <i>Command Reference</i>
Verifying compatibility of a storage system with environmental requirements	<i>Site Requirements</i>
Upgrading Data ONTAP	<i>Data ONTAP® Upgrade Guide</i>

If you need more information about...	Go to...
Migrating the cfmode, if necessary	<i>Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations</i>
Installing and configuring the HBA in your host	Your HBA vendor documentation
Your host operating system and using its features	Refer to your operating system documentation.
PowerVM	Refer to your operating system documentation. Additional documentation on using PowerVM environments is available on the IBM web site.
Veritas Storage Foundation and its features	<p>Refer to the Veritas documentation. In particular, the following documents might be helpful:</p> <p>Note: While some of these documents apply to all versions of Veritas Storage Foundation, others apply to specific versions. Make sure that the document you get is the correct version.</p> <ul style="list-style-type: none"> • Veritas Storage Foundation and High Availability Solutions - Getting Started Guide (AIX) • Veritas Storage Foundation - Installation Guide (AIX) • Veritas Storage Foundation for Oracle RAC - Installation and Configuration Guide (AIX)
General product information, including support information	The IBM NAS Interoperability Matrices Web site at www.ibm.com/systems/storage/network/interophome.html

Related information

NAS Interoperability Matrices Web site - www.ibm.com/systems/storage/network/interophome.html

FC and iSCSI Configuration Guide - <http://www.ibm.com/storage/support/nas/>

Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations - <http://www.ibm.com/storage/support/nas/>

Veritas Storage Foundation DocCentral - <http://sfdoccentral.symantec.com/>

Copyright and trademark information

Copyright information

Copyright ©1994 - 2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2011 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 retrievals 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 retrievals 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 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, FAServer, FilerView, FlexCache, FlexClone, 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), ONTAPI, OpenKey, RAID-DP, ReplicatorX, SANscreen, SecureAdmin, SecureShare, Select, Shadow Tape, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, 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, and Web Filer 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 is a licensee of the CompactFlash and CF Logo trademarks.

NetApp 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

adapter

verifying configuration 118

AIX

Host Utilities environments 14

support for non-English language versions 107

AIX Host Utilities

AIX MPIO

See MPIO

ALUA

enabling after Host Utilities installation 116

supported AIX configurations 106

supported with FC 105

verifying 116

APM

available from Symantec 51

installed with certain Veritas versions 51

installing 55

obtaining from Symantec 54

required in Veritas environments 53

uninstalling 57

used with Veritas 51

array types

displaying with sanlun 58

ASL

available from Symantec 51

determining the ASL version 54

error messages 60

installing 55

obtaining from Symantec 54

required in Veritas environments 53

uninstalling 57

used with Veritas 51

C

configurations

finding more information 123

D

Discovery Filename

adding storage system entries 48

DLPAR

supported with Host Utilities 104

documentation

finding more information 123

dotpaths

installed by Host Utilities 13

E

Environments

finding information on different environments 18

error messages

ASL 60

F

fast_fail

configuring 36

FC

using on with iSCSI 104

FC protocol

ALUA 105

ALUA configurations 106

fc_err_recov

configuring for fast_fail 36

fc_err_recov parameter

required for Veritas 51

finding more information 123

G

GPFS

supported with some configurations 105

H

HBA

removing protocol device 44

configuring 45, 47

displaying information with sanlun 79

verifying configuration 118

verifying for iSCSI 117

verifying storage system connection 118

HBAs

using with the Host Utilities 104

host settings

software package 26

Host Utilities

ASL and APM for Veritas 51

ASL and APM requirements 53
contents 13
downloading software packages 26
environments 14
FC and iSCSI on one host 104
finding more information 123
general components 13
information on the different environments 18
installing with a command line 33
installing with SMIT 28
methods for installing 25
MPIO environment 14
planning installation 19
PowerVM environment 14
preparing to uninstall 37
Release Notes 13
software packages 25
support for non-English versions 107
uncompressing software packages 26
uninstalling with a command line 41
uninstalling with SMIT 37
using VxVM 58
Veritas environment 14

I

information
finding more 123
installation
iSCSI configuration 43
downloading software packages 26
overview 19
planning 19
installing
methods for installing the Host Utilities 25
iSCSI
environments for Host Utilities 18
using on with FC 104
verifying adapter configuration 118
verifying HBA 117
verifying HBA logged in to storage system 118
iSCSI protocol
configuration 43
adding discovery information 48
changing node names 43
configuring the HBA 45, 47
configuring the software initiator 47
preparing HBA 44

J

Jumbo frames
not supported with SAN booting 107

L

languages
support for non-English versions 107
Live Partition Mobility
supported with some configurations 105
LUNs
configuration overview 63
displaying with sanlun 77

M

man pages
installed by Host Utilities 13
MPIO
software package 25
ALUA 14
ALUA configurations 106
ALUA support 105
environment for Host Utilities 14
environments for Host Utilities 18
getting Host Utilities software package 26
multipathing 14
protocols 14
MPIO/FC
environments for Host Utilities 18
multipathing
options 104
using sanlun with Veritas 58

N

node names
changing 43
non-English versions supported 107
NPIV
PowerVM 18, 70
nSANity
diagnostic utility 13
installing 120

P

polling interval
recommended values 59
PowerVM
ALUA 14
discover LUNs in NPIV environment 70
discover LUNs in VSCSI environment 66

environment for Host Utilities 14
 finding information on different environments 18
 getting Host Utilities software package 26
 multipathing 14
 NPIV 18
 protocols 14
 VIO Server 14
 VSCSI 18
 problems
 checking troubleshooting 110
 publications
 finding more information 123

Q

queue depth
 verifying 35

R

Release Notes
 checking 109
 check periodically 13
 requirements
 finding more information 123
 restore policy
 recommended value 59

S

SAN
 software package 25
 SAN booting
 advantages 107
 MPIO environments 81
 not supported on Jumbo frames 107
 SAN Toolkit
 getting Host Utilities software package 26
 SAN_Tool_Kit
 software package 25
 san_version command
 installed by Host Utilities 13
 sanlun utility
 displaying array types 58
 displaying HBA information 79
 displaying LUNs 77
 displaying multipathing for Veritas 58
 installed by Host Utilities 13
 SMIT

installing Host Utilities 28
 uninstalling Host Utilities 37
 software initiator
 configuring 47
 software packages
 downloading Host Utilities software 26
 installing Host Utilities with a command line 33
 installing Host Utilities with SMIT 28
 MPIO and PowerVM 25
 SAN_Tool_Kit 25
 uncompressing 26
 uninstalling Host Utilities with a command line 41
 uninstalling Host Utilities with SMIT 37
 Veritas SAN package 25
 storage systems
 adding to Discovery Filename 48
 Symantec
 getting the ASL and APM 54
 provides ASL, APM 51

T

troubleshooting
 checking Release Notes 109
 finding more information 123

U

uninstalling
 before you uninstall 37
 Host Utilities 37
 Host Utilities using a command line 41

V

Veritas
 ALUA support 105
 APM 51
 ASL 51
 ASL and APM requirements 53
 displaying multipathing 58
 drivers 14
 environment for Host Utilities 14
 fc_err_recov parameter 51
 getting Host Utilities software package 26
 multipathing 14
 protocols 14
 restore daemon settings 59
 SAN software package 25

setting the fast_fail option	36	PowerVM	14
setup issues	14	Virtual I/O Server/Client	14
using VxVM to display LUN paths	58	VSCSI	
Veritas DMP		PowerVM	18, 66
environments for Host Utilities	18	VxVM	
SAN booting instructions	107	displaying LUN paths	58
VIO			



NA 210-05152_A0, Printed in USA

GC27-3925-00

