

IBM System Storage N series



Solaris Host Utilities 6.1 Installation and Setup Guide

Contents

Preface	9
Supported features	9
Websites	9
Getting information, help, and service	9
Before you call	10
Using the documentation	10
Hardware service and support	10
Firmware updates	10
How to send your comments	11
Changes to this document: January 2013	12
The Solaris Host Utilities	13
What the Host Utilities contain	13
Supported Solaris environments and protocols	14
How to find instructions for your Solaris Host Utilities environment	16
Planning the installation and configuration of the Host Utilities	18
Overview of prerequisites for installing and setting up the Host Utilities	18
Host Utilities installation overview	19
iSCSI configuration	20
LUN configuration	20
(FC) Information on setting up the drivers	22
General information on getting the driver software	22
Downloading and extracting the Emulex software	22
Solaris drivers for Emulex HBAs (emlxs)	23
Installing the EMLXemlxu utilities	24
Determining Emulex firmware and FCode versions for native drivers	24
Upgrading the firmware for native drivers	24
Updating your FCode HBAs with native drivers	25
Solaris drivers for QLogic HBAs (qlc)	25
Downloading and extracting the QLogic software	25
Installing the SANsurfer CLI package	26
Determining the FCode on QLogic cards	26
Upgrading the QLogic FCode	27

The Solaris Host Utilities installation process	29
Key steps involved in setting up the Host Utilities	29
The software packages	29
Downloading the Host Utilities software	30
Installing the Solaris Host Utilities software	31
Information on upgrading or removing the Solaris Host Utilities	33
Upgrading the Solaris Host Utilities or reverting to another version	33
Methods for removing the Solaris Host Utilities	33
Uninstalling Solaris Host Utilities 6.x, 5.x, 4.x, 3.x	34
Uninstalling the Attach Kit 2.0 software	35
(iSCSI) Additional configuration for iSCSI environments	36
iSCSI node names	36
(iSCSI) Recording the initiator node name	37
(iSCSI) Storage system IP address and iSCSI static, ISNS, and dynamic discovery	37
(Veritas DMP/iSCSI) Support for iSCSI in a Veritas DMP environment	37
(iSCSI) CHAP authentication	38
(iSCSI) Configuring bidirectional CHAP	38
(iSCSI) Data ONTAP upgrades can affect CHAP configuration	39
About the host_config command	40
host_config options	40
host_config command examples	42
(Veritas DMP/FC) Tasks for completing the setup of a Veritas DMP stack	47
(Veritas DMP) Before you configure the Host Utilities for Veritas DMP	47
(Veritas DMP) sd.conf and ssd.conf variables for systems using native drivers	48
Tasks for completing the setup of a MPxIO stack	49
Before configuring system parameters on a MPxIO stack	49
Parameter values for systems using MPxIO	50
(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	52
(Veritas DMP) ASL and APM installation overview	53
(Veritas) Determining the ASL version	53

(Veritas) How to get the ASL and APM	54
(Veritas DMP) Installing the ASL and APM software	55
(Veritas DMP) Tasks to perform before you uninstall the ASL and APM ...	56
(Veritas DMP) What an ASL array type is	60
(Veritas DMP) The storage system's FC failover mode or iSCSI configuration and the array types	60
(Veritas DMP) Using VxVM to display available paths	60
(Veritas) Displaying multipathing information using sanlun	61
(Veritas DMP) Veritas environments and the fast recovery feature	61
(Veritas DMP) The Veritas DMP restore daemon requirements	62
(Veritas DMP) Setting the restore daemon interval for 5.0 MP3 and later ...	62
(Veritas DMP) Probe Idle LUN settings	63
(Veritas DMP) DMP Path Age Settings	63
(Veritas) Information about ASL error messages	64
LUN configuration and the Solaris Host Utilities	65
Overview of LUN configuration and management	65
Tasks necessary for creating and mapping LUNs	66
How the LUN type affects performance	67
Methods for creating igroups and LUNs	67
Best practices for creating igroups and LUNs	67
(iSCSI) Discovering LUNs	68
Solaris native drivers and LUNs	68
(Solaris native drivers) Getting the controller number	68
(Solaris native drivers) Discovering LUNs	69
Labeling the new LUN on a Solaris host	70
Methods for configuring volume management	72
The sanlun utility	73
Displaying host LUN information with sanlun	73
Displaying path information with sanlun	75
Explanation of the sanlun lun show -p output	75
Displaying host HBA information with sanlun	76
About the collectinfo command	78
collectinfo options	78
collectinfo command example	79
SAN boot LUNs in a Solaris Native FC environment	81
Prerequisites for creating a SAN boot LUN	81

General SAN Boot Configuration Steps	82
About SPARC OpenBoot	82
About setting up the Oracle native HBA for SAN booting	83
SPARC: Changing the Emulex HBA to SFS mode	83
SPARC: Changing the QLogic HBA to enable FCode compatibility	85
Information on creating the bootable LUN	87
Veritas DMP Systems: Gathering SAN boot LUN information	87
Native MPxIO Systems: Gathering SAN boot LUN information	90
Gathering source disk information	92
Partitioning and labeling SAN boot LUNs	93
UFS File systems: Copying data from locally booted disk	96
ZFS File systems: Copying data from locally booted disk	98
Solaris 10 ZFS: Copying data from a locally booted disk	98
Solaris 11 ZFS: Copying data from a locally booted disk	100
Making the SAN boot LUN bootable	102
SPARC: Installing the boot block	102
X64: Installing GRUB information	104
Configuring the host to boot from the SAN boot LUN	107
Configuring the host to boot from the SAN boot LUN on SPARC-based systems	107
Configuring the host to boot from the SAN boot LUN on X64-based systems	108
Veritas DMP: Enabling root encapsulation	110
Supported Solaris and Data ONTAP features	111
Features supported by the Host Utilities	111
HBAs and the Solaris Host Utilities	111
Multipathing and the Solaris Host Utilities	112
iSCSI and multipathing	112
Volume managers and the Solaris Host Utilities	112
(FC) ALUA support with certain versions of Data ONTAP	113
(FC) Solaris Host Utilities configurations that support ALUA	113
Oracle VM Server for SPARC (Logical Domains) and the Host Utilities	114
SAN booting and the Host Utilities	114
Support for non-English versions of Solaris operating systems	115
High-level look at Host Utilities Veritas DMP stack	115
High-level look at Host Utilities MPxIO stack	116

Protocols and configurations supported by the Solaris Host Utilities .	118
Notes about the supported protocols	118
The FC protocol	118
The iSCSI protocol	118
Supported configurations	119
Troubleshooting	120
About the troubleshooting sections that follow	120
Check the version of your host operating system	120
Confirm the HBA is supported	121
(MPxIO, native drivers) Ensure that MPxIO is configured correctly for	
ALUA on FC systems	122
Ensure that MPxIO is enabled on SPARC systems	123
(MPxIO) Ensure that MPxIO is enabled on iSCSI systems	123
(MPxIO) Verify that MPxIO multipathing is working	124
(Veritas DMP) Check that the ASL and APM have been installed	125
(Veritas) Check VxVM	126
(MPxIO) Check the Solaris Volume Manager	126
(MPxIO) Check settings in <code>ssd.conf</code> or <code>sd.conf</code>	126
Check the storage system setup	127
(MPxIO/FC) Check the ALUA settings on the storage system	127
Verifying that the switch is installed and configured	128
Determining whether to use switch zoning	128
Power up equipment in the correct order	128
Verify that the host and storage system can communicate	129
Possible iSCSI issues	129
(iSCSI) Verify the type of discovery being used	129
(iSCSI) Bidirectional CHAP does not work	129
(iSCSI) LUNs are not visible on the host	129
Possible MPxIO issues	131
(MPxIO) <code>sanlun</code> does not show all adapters	131
(MPxIO) Solaris log message says data not standards compliant	131
Installing the nSANity data collection program	131
LUN types, OS label, and OS version combinations for achieving	
aligned LUNs	133
Where to find more information	135
Copyright information	137

Trademark information	138
Index	141

Preface

Supported features

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

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Websites

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

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

Getting information, help, and service

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

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

Before you call

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

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 9) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 9).

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

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

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

Hardware service and support

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

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 9).

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

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

How to send your comments

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

Be sure to include the following:

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

Changes to this document: January 2013

Several changes have been made to this document since it was published for the Solaris Host Utilities 6.0 release.

This document has been updated to add the following information:

- The section on SAN boot LUNs in a Solaris Veritas DMP environment with FC has been updated to support changes to SAN boot.

The Solaris Host Utilities

The Solaris Host Utilities are a collection of components that enable you to connect Solaris hosts to N series storage systems running Data ONTAP.

Once connected, you can set up logical storage units known as LUNs (Logical Unit Numbers) on the storage system.

Note: Previous versions of the Host Utilities were called FCP Solaris Attach Kits and iSCSI Support Kits.

The following sections provide an overview of the Solaris Host Utilities environments, and information on what components the Host Utilities supply.

What the Host Utilities contain

The Host Utilities bundle numerous software tools into a SAN Toolkit.

Note: This toolkit is common across all the Host configurations: FCP and iSCSI with MPxIO and Veritas DMP. As a result, some of its contents apply to one configuration, but not another. Having a program or file that does not apply to your configuration does not affect performance.

The toolkit contains the following components:

- `san_version` command. This command displays the version of the SAN Toolkit that you are running.
- `sanlun` utility. This utility displays information about LUNs on the storage system that are available to this host.
- `host_config` command. This command modifies the SCSI retry and timeout values in the following files:
 - `/kernel/drv/sd.conf`
 - `/kernel/drv/ssd.conf`

It also adds or deletes the `symmetric-option` and `IBM VIP/PID` in the `/kernel/drv/scsi_vhci.conf` file.

- `collectinfo` command. This optional command collects anonymous support information about system hosts such as HBA types, volume manager configuration and operating system. When the command is executed, this information is pushed to the IBM N series controller. If the IBM N series controller has AutoSupport enabled, this information is part of the payload that is used to ensure future releases of the Host Utilities meet customer needs. This command must be initiated by the user with proper controller login credentials.
- The man pages for `sanlun` and the diagnostic utilities.

Note: Previous versions of the Host Utilities also included diagnostics programs. These programs have been replaced by the nSANity Diagnostic and Configuration Data Collector and are no longer installed with the Host Utilities. The nSANity program is not part of the Host Utilities. You should download, install, and execute nSANity 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 Changed by the Host Utilities*
- *Quick Command Reference*

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

Supported Solaris environments and protocols

The Host Utilities support several Solaris environments.

For details on which environments are supported, see the online IBM Interoperability Matrix.

The following table summarizes key aspects of the two main environments.

Solaris Environment	Notes
Veritas DMP	<ul style="list-style-type: none"> • This environment uses Veritas Storage Foundation and its features. • Multipathing: Veritas Dynamic Multipathing (DMP) with either Solaris native drivers (Leadville) or iSCSI. • Volume management: Veritas Volume Manager (VxVM). • Protocols: Fibre Channel (FC) and iSCSI. • Software package: Install the software packages in the compressed download file for your host platform. • Setup issues: <ul style="list-style-type: none"> • You might need to perform some driver setup. • The Symantec Array Support Library (ASL) and Array Policy Module (APM) might need to be installed. See the N series support website (accessed and navigated as described in Websites on page 9) for the most current information on system requirements. • Configuration issues: <ul style="list-style-type: none"> • SPARC systems using FCP require changes to the parameters in the <code>/kernel/drv/ssd.conf</code> file. • SPARC systems using iSCSI require changes to the parameters in the <code>/kernel/drv/sd.conf</code> file. • All x86 systems require changes to the parameters in the <code>/kernel/drv/sd.conf</code> file. <p>Note: Asymmetric Logical Unit Access (ALUA) is supported with Veritas 5.1 and later.</p>

Solaris Environment	Notes
MPxIO (Native MultiPathing)	<ul style="list-style-type: none"> • This environment works with features provided by the Solaris operating system. It uses Solaris StorEdge SAN Foundation Software. • Multipathing: Solaris StorageTek Traffic Manager (MPxIO) or the Solaris iSCSI Software Initiator. • Volume management: Solaris Volume Manager (SVM), ZFS, or VxVM. • Protocols: FC and iSCSI. • Before Data ONTAP 8.1.1, ALUA is only supported in FC environments. (It is also supported with one older version of the iSCSI Support Kit: 3.0.). Data ONTAP 8.1.1 supports ALUA in FC and iSCSI environments. • Software package: Download the compressed file associated with your system's processor (SPARC or x86/64) and install the software packages in that file. • Setup issues: None. • Configuration issues: <ul style="list-style-type: none"> • Systems using SPARC processors require changes to the parameters in the <code>/kernel/drv/ssd.conf</code> file. • Systems using x86 processors require changes to the parameters in the <code>/kernel/drv/sd.conf</code> file.

Related information

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

How to find instructions for your Solaris 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 "**PowerVM**," 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 Solaris Host Utilities environments, there is no qualifier in the title.

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

Qualifier	The section that follows applies to
(Veritas DMP)	Environments using Veritas DMP as the multipathing solution.
(Veritas DMP/native)	Veritas DMP environments that use Solaris native drivers.
(Veritas DMP/iSCSI)	Veritas DMP environments that use the iSCSI protocol.
(MPxIO)	Environments using MPxIO as the multipathing solution. Currently, all MPxIO environments use native drivers.
(MPxIO/FC)	MPxIO environments using the FC protocol.
MPxIO/iSCSI	MPxIO environments using the iSCSI protocol.
(FC)	Environments using the Fibre Channel protocol. Note: Unless otherwise specified, FC refers to both FC and FCoE in this guide.
(iSCSI)	Environments using the iSCSI protocol.

There is also information about using the Host Utilities in a Solaris environment in the *Release Notes* and the Solaris Host Utilities reference documentation. You can download all the Host Utilities documentation from the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Planning the installation and configuration of the Host Utilities

Installing the Host Utilities and setting up your system involves a number of tasks that are performed on both the storage system and the host.

You should plan your installation and configuration before you install the Host Utilities. The following sections help you do this by providing a high-level look at the different tasks you need to perform to complete the installation and configuration of the Host Utilities. The detailed steps for each of these tasks are provided in the chapters that follow these overviews.

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

Overview of prerequisites for installing and setting up the Host Utilities

As you plan your installation, keep in mind that there are several tasks that you should perform before you install the Host Utilities.

The following is a summary of the tasks you should perform before installing the Host Utilities:

1. Verify your system setup:

- Host operating system and appropriate updates
- HBAs or software initiators
- Drivers
- **Veritas environments only:** Veritas Storage Foundation, the Array Support Library (ASL) for the storage controllers, and if you are using Veritas Storage Foundation 5.0, the Array Policy Module (APM)

Note: Make sure you have the Veritas Volume Manager (VxVM) installed before you install the ASL and APM software. The ASL and APM are available from the Symantec Website.

- Volume management and multipathing, if used.
- Storage system with Data ONTAP installed.
- **iSCSI environments only:** Record or set the host's iSCSI node name.
- **FC environments only:** Switches, if used.

Note: For information about supported topologies, see the *SAN Configuration Guide (called Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier)* for your version of Data ONTAP, which is available online.

For the most current information about system requirements, see the N series support website (accessed and navigated as described in [Websites](#) on page 9).

2. Verify that your storage system is:
 - Licensed correctly for the protocol you are using and running that protocol service.
 - **For Data ONTAP operating in 7-Mode only:** Using the recommended cfmode (single-image).
 - Configured to work with the target HBAs, as needed by your protocol.
 - Set up to work with the Solaris host and the initiator HBAs or software initiators, as needed by your protocol.
 - **FC active/active environments only:** Set up to work with ALUA, if it is supported by your multipathing solution.

Note: For Data ONTAP operating in 7-Mode, ALUA is not supported with iSCSI.
 - Set up with working volumes and qtrees (if desired).
3. **FC environments only:** If you are using a switch, verify that it is:
 - Set up correctly
 - Zoned
 - Cabled correctly
 - Powered on in the correct order: switches, disk shelves, storage systems, and then the host
4. Confirm that the host and the storage system can communicate.
5. If you currently have the Host Utilities installed, remove that software.

Related information

N series support website - www.ibm.com/storage/support/nseries/

Host Utilities installation overview

The actual installation of the Host Utilities is fairly simple. As you plan the installation, you need to consider the tasks you must perform to get the Host Utilities installed and set up for your environment.

The following is a high-level overview of the tasks required to install the Host Utilities. The chapters that follow provide details on performing these tasks.

1. Get a copy of the compressed Host Utilities file, which contains the software package for your multipathing solution and the SAN Toolkit software package.
 - Download the compressed file containing the packages for your multipathing solution.
 - Extract the software packages from the compressed file that you downloaded.

2. Install the Host Utilities software packages. You must be logged in as `root` to install the software.
 - From the directory containing the extracted software packages, use the `pkgadd -d` command to install the Host Utilities package for your stack.
 - Set the driver and system parameters. You do this using the `host_config` command.

Note: You can also set the parameters manually.
3. Complete the configuration based on your environment.
 - **(iSCSI)** There are several tasks you need to perform to get your iSCSI environment set up. They include recording the iSCSI node name, setting up the initiator, and, optionally, setting up CHAP.
 - **(Veritas)** Make sure you have the ASL and APM correctly installed and set up if required for your Veritas version. See the N series support website (accessed and navigated as described in [Websites](#) on page 9) for the most current information on system requirements.

iSCSI configuration

If you are using the iSCSI protocol, then you must perform some additional configuration to set it up correctly for your environment.

1. Record the host's iSCSI node name.
2. Configure the initiator with the IP address for each storage system. You can use static, ISNS, or sendtargets.
3. **Veritas iSCSI environment only:** Make sure MPxIO is disabled. If you had an earlier version of the Host Utilities installed, you might need to remove the MPxIO settings that it set up and then reboot your host. To remove these settings, do one of the following:
 - Use the `host_config` command to remove both the IBM VID/PID and the symmetric-option from the `/kernel/drv/scsi_vhci.conf` file for Solaris 10 or `/etc/driver/drv/scsi_vhci.conf` file for Solaris 11.
 - Manually edit the `/kernel/drv/scsi_vhci.conf` for Solaris 10 or `/etc/driver/drv/scsi_vhci.conf` file for Solaris 11 and remove the VID/PID entries.
4. (Optional) Configure CHAP on the host and the storage system.
5. If you are running Solaris 10u9 or later in conjunction with Data ONTAP 8.1 or later, adjust the value of `conn-login-max` to 60 on the client iSCSI initiator using the following command:

```
# iscsiadm modify initiator-node -T conn-login-max=60
```

LUN configuration

To complete your setup of the Host Utilities, you need to create LUNs and get the host to see them.

Configure the LUNs by performing the following tasks:

- Create at least one igroup and at least one LUN and map the LUN to the igroup.
One way to create igroups and LUNs is to use the `lun setup` command. Specify `solaris` as the value for the `ostype` attribute. You will need to supply a WWPN for each of the host's HBAs or software initiators.
- **MPxIO FC environments only:** Enable ALUA, if you have not already done so.
- Configure the host to discover the LUNs.
 - **Native drivers:** Use the `/usr/sbin/cfgadm -c configure cx` command, where `x` is the controller number of the HBA where the LUN is expected to be visible.
- Label the LUNs using the Solaris format utility (`/usr/sbin/format`).
- Configure the volume management software.
- Display information about the LUNs and HBA.
You can use the `sanlun` command to do this.

(FC) Information on setting up the drivers

For Emulex-branded HBAs, the Emulex Utilities are required to update the firmware and boot code. These utilities can be downloaded directly from Emulex.

General information on getting the driver software

You can get the driver software from the company website for your HBA.

To determine which drivers are supported with the Host Utilities, check the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

- Emulex HBAs with Solaris native drivers: The Emulex software, including the Emulex utility programs and documentation, is available from the Solaris OS download section on the Emulex site.
- QLogic-branded HBAs: The QLogic SANsurfer CLI software and documentation are available on the QLogic support site. QLogic provides a link to its IBM partner sites. You only need this software if you have to manipulate the FCode versions on QLogic-branded HBAs for SAN booting.
- Oracle-branded HBAs: You can also use certain Oracle-branded HBAs. For more information on working with them, see the patch Readme file that Oracle provides.

Related information

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

[Emulex partner site](#)

[QLogic partner site](#)

Downloading and extracting the Emulex software

The following steps tell you how to download and extract the Emulex software and firmware.

About this task

If your HBA uses an earlier version of the firmware than is supported by the Host Utilities, you need to download new firmware when you download the rest of the Emulex software. To determine which firmware versions are supported, check the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

Steps

1. On the Solaris host, create a download directory for the Emulex software, such as: `mkdir /tmp/emulex`, and change to that directory.
2. To download the Emulex driver and firmware, go to the location on the Emulex Website for the type of drivers you are using:
 - For Emulex HBAs using Solaris native drivers, go to the Solaris OS download section.
3. Follow the instructions on the download page to get the driver software and place it in the `/tmp/emulex` directory you created.
4. Use the `tar xvf` command to extract the software from the files you downloaded.

Note: If you are using Emulex Utilities for Solaris native drivers, the .tar file you download contains two additional .tar files, each of which contains other .tar files. The file that contains the EMLXemlxu package for native drivers is `emlxu_kit-<version>-sparc.tar`.

The following command line show how to extract the software from the files for the Emulex Utility bundle for use with Solaris Native drivers:

```
tar xvf solaris-HBAnyware_version-utility_version-subversion.tar
```

Related information

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

Solaris drivers for Emulex HBAs (emlxs)

The Host Utilities supports Emulex HBAs with Solaris native drivers. The Emulex software for these drivers is provided as .tar files. You need the .tar files containing the Emulex Fibre Channel Adapter (FCA) Utilities (EMLXemlxu).

The FCA utilities manage the firmware and FCode of the Emulex HBAs with Solaris native drivers. To install and use these utilities, follow the instructions in the Emulex *FCA Utilities Reference Manual*.

The sections that follow contain information on what you need to do to set up these drivers for the Host Utilities' Veritas DMP environment.

Installing the EMLXemlxu utilities

After you extract the EMLXemlxu utilities, you must install the EMLXemlxu package.

Step

1. Run the `emlxu_install` command to install the EMLXemlxu package:

```
# ./emlxu_install
```

Note: For more information on installing and using these utilities, see the Emulex *FCA Utilities Reference Manual*.

Determining Emulex firmware and FCode versions for native drivers

Make sure you are using the Emulex firmware recommended for the Host Utilities when using Emulex-branded HBAs.

About this task

To determine which version of firmware you should be using and which version you are actually using, complete the following steps:

Steps

1. Check the N series support website (accessed and navigated as described in [Websites](#) on page 9) to determine the current firmware requirements.

2. Run the `emlxadm` utility. Enter:

```
/opt/EMLXemlxu/bin/emlxadm
```

The software displays a list of available adapters.

3. Select the device that you want to check.

The software displays a menu of options.

4. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt.

Upgrading the firmware for native drivers

If you are not using the Emulex firmware recommended for the Host Utilities using native drivers, you must upgrade your firmware.

About this task

Note: Oracle-branded HBAs have the proper firmware version pushed to the card by the native driver.

Steps

1. Run the `emlxadm` utility. Enter:

```
/opt/EMLXemlxu/bin/emlxadm
```

The software displays a list of available adapters.

2. At the `emlxadm>` prompt, enter:

```
download_fw filename
```

The firmware is loaded onto the selected adapter.

3. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt.
4. Reboot your host.

Updating your FCode HBAs with native drivers

If you are not using the correct FCode for HBAs using native drivers, you must upgrade it.

Steps

1. Run the `emlxadm` utility. Enter:

```
/opt/EMLXemlxu/bin/emlxadm
```

The software displays a list of available adapters.

2. Select the device you want to check.

The software displays a menu of options.

3. At the `emlxadm>` prompt, enter:

```
download_fcode filename
```

The FCode is loaded onto the selected adapter.

4. Exit the `emlxadm` utility by entering `q` at the `emlxadm>` prompt

Solaris drivers for QLogic HBAs (qlc)

The Host Utilities support QLogic-branded and Oracle-branded QLogic OEM HBAs that use the native driver (qlc) software. The following sections provide information on setting up these drivers.

Downloading and extracting the QLogic software

If you are using QLogic drivers, you must download and extract the QLogic software and firmware.

Steps

1. On the Solaris host, create a download directory for the QLogic software. Enter:

```
mkdir /tmp/qlogic
```

2. To download the SANsurfer CLI software, go to the [QLogic website](#) and click the Downloads link.
3. Under “OEM Models,” click IBM.
4. Click the link for your card type.
5. Choose the latest multiflash or bios image available and save it to the `/tmp/qlogic` directory on your host
6. Change to the `/tmp/qlogic` directory and uncompress files that contain the SANsurfer CLI software package. Enter:

```
uncompress scli-version.SPARC-X86.Solaris.pkg.Z
```

Installing the SANsurfer CLI package

After you extract the QLogic files, you need to install the SANsurfer CLI package.

Steps

1. Install the SANsurfer CLI package using the `pkgadd` command. Enter:

```
pkgadd -d /tmp/qlogic/scli-version.SPARC-X86.Solaris.pkg
```

2. From the directory where you extracted the QLogic software, unzip the FCode package. Enter:

```
unzip fcode_filename.zip
```
3. For instructions about updating the FCode, please see the “Upgrading the QLogic FCode.”

Related tasks

[Upgrading the QLogic FCode](#) on page 27

Determining the FCode on QLogic cards

If you are not using the FCode recommended for the Host Utilities, you must upgrade it.

Steps

1. Check the N series support website (accessed and navigated as described in [Websites](#) on page 9) to determine the current FCode requirements.
2. Run the `scli` utility to determine whether your FCode is current or needs updating. Enter:

```
/usr/sbin/scli
```

The software displays a menu.

3. Select option 3 (HBA Information Menu).

The software displays the HBA Information Menu.

4. Select option 1 (Information).
The software displays a list of available ports.
5. Select the adapter port for which you want information.
The software displays information about that HBA port.
6. Write down the FCode version and press Return.
The software displays a list of available ports.
7. Repeat steps 5 and 6 for each adapter you want to query. When you have finished, select option 0 to return to the main menu.
The software displays the main menu.
8. To exit the scli utility, select option 13 (Quit).

Upgrading the QLogic FCode

If you are not using the correct FCode for HBAs using QLogic, you must upgrade it.

Steps

1. Run the scli utility. Enter:

```
/usr/sbin/scli
```


The software displays a menu.
2. Select option 8 (HBA Utilities).
The software displays a menu.
3. Select option 3 (Save Flash).
The software displays a list of available adapters.
4. Select the number of the adapter for which you want information.
The software displays a file name to use.
5. Enter the name of the file into which you want to save the flash contents.
The software backs up the flash contents and then waits for you to press Return.
6. Press Return.
The software displays a list of available adapters.
7. If you are upgrading more than one adapter, repeat steps 4 through 6 for each adapter.
8. When you have finished upgrading the adapters, select option 0 to return to the main menu.
9. Select option 8 (HBA Utilities).
The software displays a menu.

10. Select option 1 (Update Flash).

The software displays a menu of update options.

11. Select option 1 (Select an HBA Port)

The software displays a list of available adapters.

12. Select the appropriate adapter number.

The software displays a list of Update ROM options.

13. Select option 1 (Update Option ROM).

The software requests a file name to use.

14. Enter the file name of the multiflash firmware bundle that you extracted from the file you downloaded from QLogic. The file name should be similar to q24mf129.bin

The software upgrades the FCode.

15. Press Return.

The software displays a menu of update options.

16. If you are upgrading more than one adapter, repeat steps 11 through 15 for each adapter.

17. When you have finished, select option 0 to return to the main menu.

18. To exit the scli utility, select option 13 (Quit).

The Solaris Host Utilities installation process

The Solaris Host Utilities installation process involves several tasks. You must make sure your system is ready for the Host Utilities, download the correct copy of the Host Utilities installation file, and install the software. The following sections provide information on tasks making up this process.

Key steps involved in setting up the Host Utilities

Setting up the Host Utilities on your system involves both installing the software package for your stack and then performing certain configuration steps based on your stack.

Before you install the software, confirm the following:

- Your host system meets requirements and is set up correctly. Check the Interoperability Matrix to determine the current hardware and software requirements for the Host Utilities.
- **(Veritas DMP)** If you are using a Veritas environment, make sure Veritas is set up. For some Veritas versions, you will need to install the Symantec Array Support Library (ASL) and Array Policy Module (APM) for ONTAP storage systems. See the online IBM Interoperability Matrix for specific system requirements.
- You do not currently have a version of the Solaris Host Utilities, Solaris Attach Kit, or the iSCSI Support kit installed. If you previously installed one of these kits, you must remove it before installing a new kit.
- You have a copy of the Host Utilities software.

When you have installed the software, you can use the `host_config` script it provides to complete your setup and configure your host parameters.

After you install the Host Utilities software, you will need to configure the host system parameters. The configuration steps you perform depend on which environment you are using:

- Veritas DMP
- MPxIO

In addition, if you are using the iSCSI protocol, you must perform some additional setup steps.

The software packages

There are two Host Utilities software distribution packages.

You only need to install the file that is appropriate for your system. The two packages are:

- **SPARC processor systems:** Install this software package if you have either a Veritas DMP environment or an MPxIO environment that is using a SPARC processor.

- **x86/64 systems:** Install this software package if you have either a Veritas environment or an MPxIO environment that is using an x86/64 processor.

Downloading the Host Utilities software

You can download the Host Utilities software package for your environment or you can install it from the physical media.

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 N series support website (accessed and navigated as described in *Websites* on page 9).

Both the FC protocol and the iSCSI protocol use the same version of the Host Utilities software.

Step

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 N series support website (accessed and navigated as described in *Websites* on page 9).

After you finish

Next you need to uncompress the software file and then install the software using a command such as `pkgadd` to add the software to your host.

Related information

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

Installing the Solaris Host Utilities software

Installing the Host Utilities involves uncompressing the files and adding the correct software package to your host.

Before you begin

Make sure you have downloaded the compressed file containing the software package for the Host Utilities or gotten it from the physical media.

In addition, it is a good practice to check the Solaris 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.

Steps

1. Log in to the host system as root.
2. Place the compressed file for your processor in a directory on your host and go to that directory.

At the time this documentation was prepared, the compressed files were called:

- SPARC CPU: `ibm_solaris_host_utilities_6_1_sparc.tar.gz`
- x86/x64 CPU: `ibm_solaris_host_utilities_6_1_amd.tar.gz`

Note: The actual file names for the Host Utilities software might be slightly different from the ones shown in these steps. These are provided as examples of what the filenames look like, and to use in the examples that follow. The files you download are correct.

If you are installing the `ibm_solaris_host_utilities_6_1_sparc.tar.gz` file on a SPARC system, you might put it in the `/tmp` directory on your Solaris host.

The following example places the file in the `/tmp` directory and then moves to that directory:

```
# cp ibm_solaris_host_utilities_6_1_sparc.tar.gz /tmp
# cd /tmp
```

3. Unzip the file using the `gunzip` command.

The software unzips the `tar.gz` files.

The following example unzips files for a SPARC system:

```
# gunzip ibm_solaris_host_utilities_6_1_sparc.tar.gz
```

4. Untar the file. You can use the `tar xvf` command to do this.

The Host Utilities scripts are extracted to the default directory.

The following example uses the `tar xvf` command to extract the Solaris installation package for a SPARC system:

```
# tar xvf ibm_solaris_host_utilities_6_1_sparc.tar
```

5. Add the packages that you extracted from tar file to your host. You can use the `pkgadd` command to do this.

The packages are added to the `/opt/ontap/SANToolkit/bin` directory.

The following example uses the `pkgadd` command to install the Solaris installation package:

```
# pkgadd -d ./ontapSANTool.pkg
```

6. Confirm that the toolkit was successfully installed by using the `pkginfo` command or the `ls -al` command.

```
# ls -alR /opt/NTAP/SANToolkit
/opt/NTAP/SANToolkit:
total 598
drwxr-xr-x  3 root    sys          512 May  9 12:26 ./
drwxr-xr-x  3 root    sys          512 May  9 12:26 ../
-r-xr-xr-x  1 root    sys       292220 Jan  6 13:02 NOTICES.PDF*
drwxr-xr-x  2 root    sys          512 May  9 12:26 bin/

/opt/NTAP/SANToolkit/bin:
total 16520
drwxr-xr-x  2 root    sys          512 May  9 12:26 ./
drwxr-xr-x  3 root    sys          512 May  9 12:26 ../
-r-xr-xr-x  1 root    sys     4724280 May  8 23:36 collectinfo*
-r-xr-xr-x  1 root    sys     2086688 May  8 23:37 host_config*
-r-xr-xr-x  1 root    sys         995 May  8 23:36 san_version*
-r-xr-xr-x  1 root    sys     1606568 May  8 23:37 sanlun*
-r-xr-xr-x  1 root    sys         677 May  8 23:36 vidpid.dat*

# (cd /usr/share/man/man1; ls -al collectinfo.1 host_config.1 sanlun.
1)
-r-xr-xr-x  1 root    sys         6367 May  8 23:36 collectinfo.1*
-r-xr-xr-x  1 root    sys         9424 May  8 23:36 host_config.1*
-r-xr-xr-x  1 root    sys         9044 May  8 23:36 sanlun.1*
```

After you finish

To complete the installation, you must configure the host parameters for your environment:

- Veritas DMP
- MPxIO

If you are using iSCSI, you must also configure the initiator on the host.

Information on upgrading or removing the Solaris Host Utilities

You can easily upgrade the Solaris Host Utilities to a new version or remove an older version. If you are removing the Host Utilities, the steps you perform vary based on the version of the Host Utilities or Attach Kit that is currently installed. The following sections provide information on upgrading and removing the Host Utilities.

Upgrading the Solaris Host Utilities or reverting to another version

You can upgrade to a newer version of the Host Utilities or revert to a previous version without any effect on system I/O.

Steps

1. Use the Solaris `pkgrm` command to remove the Host Utilities software package you no longer need.

Note: Removing the software package does not remove or change the system parameter settings for that I/O stack. To remove the settings you added when you configured the Host Utilities, you must perform additional steps. You do not need to remove the settings if you are upgrading the Host Utilities.

2. Use the Solaris `pkgadd` command to add the appropriate Host Utilities software package.

Methods for removing the Solaris Host Utilities

There are two standard methods for uninstalling the Host Utilities or Attach Kit from your system. The method you use depends on the version of the kit that is installed.

- For Solaris Host Utilities 6.x, 5.x, 4.x, or 3.x, use the `pkgrm` command to remove the software package.
- For Solaris Attach Kit 2.0, use the `uninstall` script included with the Attach Kit to uninstall the software package.

Uninstalling Solaris Host Utilities 6.x, 5.x, 4.x, 3.x

If you have the Solaris Host Utilities 6.x, 5.x, 4.x, or 3.0 installed, you can use the `pkgrm` command to remove the software. If you want to revert to the saved parameter values, you must perform additional steps.

Steps

1. If you want to remove the parameters that were set when you ran the `host_config` command or that you set manually after installing the Host Utilities and restore the previous values, you can do one of the following:
 - Replace the system files with the backup files you made before changing the values.
 - **(Sun native drivers)** SPARC systems and systems: `/kernel/drv/ssd.conf`.
 - **(Sun native drivers)** x86/64 systems: `/kernel/drv/sd.conf`.
 - **(Veritas DMP)** Replace `/kernel/drv/sd.conf`
 - Use the `host_config -cleanup` command to revert to the saved values.

Note: You can only do this once.

2. Use the `pkgrm` command to remove the Solaris Host Utilities software from the `/opt/ontap/SANToolkit/bin` directory.

The following command line removes the Host Utilities software package.

```
# pkgrm ontapSANTool
```

3. You can disable MPxIO by using `stmsboot`:

- **(For FCP):**

```
# /usr/sbin/stmsboot -D fp -d
```

Answer "n" when prompted to reboot your host.

- **(For iSCSI):**

```
# /usr/sbin/stmsboot -D iscsi -d
```

Answer "n" when prompted to reboot your host.

4. To enable the changes, reboot your system using the following commands:

```
# touch /reconfigure
# init 6
```

Uninstalling the Attach Kit 2.0 software

If you have the Solaris Attach Kit 2.0 installed, complete the following steps to remove the software.

Steps

1. Ensure that you are logged in as root.
2. Locate the Solaris Attach Kit 2.0 software. By default, the Solaris Attach Kit is installed in `/opt/ontapsanlun/bin`.
3. From the `/opt/ontapsanlun/bin` directory, enter the `./uninstall` command to remove the existing software.

You can use the following command to uninstall the existing software.

```
# ./uninstall
```

Note: The uninstall script automatically creates a backup copy of the `/kernel/drv/lpfc.conf` and `sd.conf` files as part of the uninstall procedure. It is a good practice, though, to create a separate backup copy before you begin the uninstall.

4. At the prompt “Are you sure you want to uninstall lpfc and sanlun packages?” enter `y`.

The uninstall script creates a backup copy of the `/kernel/drv/lpfc.conf` and `sd.conf` files to `/usr/tmp` and names them:

- `lpfc.conf.save`
- `sd.conf.save`

If a backup copy already exists, the install script prompts you to overwrite the backup copy.

5. Reboot your system.

You can use the following commands to reboot your system.

```
# touch /reconfigure  
# init 6
```

(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.

You must:

- Record the host's initiator node name. You need this information to set up your storage.
- Configure the initiator with the IP address for each storage system using either static, ISNS, or dynamic discovery.
- (Optionally) configure CHAP.

The following sections explain how to perform these tasks.

iSCSI node names

To perform certain tasks, you need to know the iSCSI node name.

Each iSCSI entity on a network has a unique iSCSI node name. This is a logical name that is not linked to an IP address.

Only initiators (hosts) and targets (storage systems) are iSCSI entities. Switches, routers, and ports are TCP/IP devices only and do not have iSCSI node names.

The Solaris software initiator uses the `iqn`-type node name format:

`iqn.yyyy-mm.backward_naming_authority:unique_device_name`

- `yyyy` is the year and `mm` is the month in which the naming authority acquired the domain name.
- `backward_naming_authority` is the reverse domain name of the entity responsible for naming this device. An example reverse domain name is `com.ibm`.
- `unique_device_name` is a free-format unique name for this device assigned by the naming authority.

The following example shows a default iSCSI node name for a Solaris software initiator:

```
iqn.1986-03.com.sun:01:0003ba0da329.43d53e48
```

(iSCSI) Recording the initiator node name

You need to get and record the host's initiator node name. You use this node name when you configure the storage system.

Steps

1. On the Solaris host console, enter the following command:

```
iscsiadm list initiator-node
```

The system displays the iSCSI node name, alias, and session parameters.

2. Record the node name for use when configuring the storage system.

(iSCSI) Storage system IP address and iSCSI static, ISNS, and dynamic discovery

The iSCSI software initiator needs to be configured with one IP address for each storage system. You can use static, ISNS, or dynamic discovery.

When you enable dynamic discovery, the host uses the iSCSI `SendTargets` command to discover all of the available interfaces on a storage system. Be sure to use the IP address of an interface that is enabled for iSCSI traffic.

Note: See the Solaris Host Utilities Release Notes for issues with regard to using dynamic discovery.

Follow the instructions in the *Solaris System Administration Guide: Devices and File Systems* to configure and enable iSCSI `SendTargets` discovery. You can also refer to the `iscsiadm` man page on the Solaris host.

(Veritas DMP/iSCSI) Support for iSCSI in a Veritas DMP environment

The Host Utilities support iSCSI with certain versions of Veritas DMP.

Check the Interoperability Matrix to determine whether your version of Veritas DMP supports iSCSI.

To use iSCSI with Veritas DMP, make sure that MPxIO is disabled. If you previously ran the Host Utilities on the host, you might need to remove the MPxIO settings in order to allow Veritas DMP to provide multipathing support.

Related information

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

(iSCSI) CHAP authentication

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

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

(iSCSI) Configuring bidirectional CHAP

Configuring bidirectional CHAP involves several steps.

About this task

For bidirectional CHAP, the target CHAP secret value you configure on the Solaris host must be the same as the outpassword value you configured on the storage system. The target CHAP username must be set to the target's iSCSI node name on the storage system. You cannot configure the target CHAP username value on the Solaris host.

Note: Make sure you use different passwords for the inpassword value and the outpassword value.

Steps

1. Set the username for the initiator.

```
iscsiadm modify initiator-node --CHAP-name sunhostname
```

2. Set the initiator password. This password must be at least 12 characters and cannot exceed 16 characters.

```
iscsiadm modify initiator-node --CHAP-secret
```

3. Tell the initiator to use CHAP authentication.

```
iscsiadm modify initiator-node -a chap
```

4. Configure bidirectional authentication for the target.

```
iscsiadm modify target-param -B enable targetIQN
```

5. Set the target username.

```
iscsiadm modify target-param --CHAP-name filerhostname targetIQN
```

6. Set the target password. Do not use the same password as the one you supplied for the initiator password. This password must be at least 12 characters and cannot exceed 16 characters.

```
iscsiadm modify target-param --CHAP-secret targetIQN
```

7. Tell the target to use CHAP authentication.

```
iscsiadm modify target-param -a chap targetIQN
```

8. Configure security on the storage system.

```
iscsi security add -i initiatorIQN -s CHAP -p initpassword -n  
sunhostname -o targetpassword -m filerhostname"
```

(iSCSI) Data ONTAP upgrades can 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.

About the `host_config` command

The `host_config` command enables you to configure your system and automatically set recommended system values. You can use the same options for the `host_config` command across all the environments supported by the Host Utilities.

The `host_config` command has the following format:

```
host_config <-setup> <-protocol fcp|iscsi|mixed> <-multipath mpzio|dmp|
non> [-noalua]
```

Note: The `host_config` command replaces the `basic_config` command, which was used with the versions of the Host Utilities before 6.0.

This command replaces the `basic_config` command and the `basic_config` command options used before 6.0.

You must be logged on as root to run the `host_config` command. The `host_config` command does the following:

- Makes setting changes for the Fibre Channel and SCSI drivers for both X86 and SPARC systems
- Provides SCSI timeout settings for both the MPxIO and DMP configurations
- Sets the VID/PID information
- Enables or disables ALUA
- Configures the ALUA settings used by MPxIO and the SCSI drivers for both X86 and SPARC systems.

Note: iSCSI is not supported with ALUA if you are running Data ONTAP operating in 7-Mode or Data ONTAP operating in Cluster-Mode before release 8.1.1.

host_config options

The `host_config` command has several options you can use. These options apply to all environments. This command is executed on the host.

Option	Description
<code>-setup</code>	Automatically sets the recommended parameters.
<code>-protocol fcp iscsi mixed</code>	Lets you specify the protocol you will be using. Enter <code>fcp</code> if you are using the FC protocol. Enter <code>iscsi</code> if you are using the iSCSI protocol. Enter <code>mixed</code> if you are using both the FC and iSCSI protocols.

Option	Description
<code>-multipath mpxio dmp none</code>	Lets you specify your multipathing environment. If you are not using multipathing, enter the argument <code>none</code> .
<code>-noalua</code>	Disables ALUA.
<code>-cleanup</code>	Deletes parameters that have been previously set and reinitializes parameters back to the OS defaults.
<code>-help -H -?</code>	Displays a list of available commands.
<code>-version</code>	Displays the current version of the Host Utilities.

Valid `host_config -setup` combinations for Data ONTAP operating in Cluster-Mode

The following parameter combinations can be used with the `host_config` command when your storage system is running Data ONTAP operating in Cluster-Mode.

- `host_config -setup -protocol fcp -multipath mpxio`
- `host_config -setup -protocol fcp -multipath dmp`
- `host_config -setup -protocol iscsi -multipath mpxio`
- `host_config -setup -protocol iscsi -multipath dmp`
- `host_config -setup -protocol mixed -multipath mpxio`
- `host_config -setup -protocol mixed -multipath dmp`

Valid `host_config -setup` combinations for Data ONTAP operating in 7-Mode

The following parameter combinations can be used with the `host_config` command when your storage system is running Data ONTAP operating in 7-Mode.

- `host_config -setup -protocol fcp -multipath mpxio`
- `host_config -setup -protocol fcp -multipath dmp`
- `host_config -setup -protocol fcp -multipath dmp -noalua`
- `host_config -setup -protocol fcp -multipath none -noalua`
- `host_config -setup -protocol iscsi -multipath mpxio -noalua`
- `host_config -setup -protocol iscsi -multipath dmp -noalua`
- `host_config -setup -protocol iscsi -multipath none -noalua`

host_config command examples

The following examples step you through the process of using the `host_config` command to configure your system.

Note: If you need to remove these changes, run the `host_config <-cleanup>` command.

Native FCP Driver with MPxIO Usage (SPARC) - Solaris 10

```
# host_config -setup -protocol fcp -multipath mpzio

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####

ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:64, throttle-min:8";

#####
#####
The following lines will be REMOVED from the /kernel/drv/
scsi_vhci.conf file
#####
#####

device-type-scsi-options-list =
"NETAPP LUN", "symmetric-option";
symmetric-option = 0x1000000;

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D fp -e      (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

Native FCP Drive with MPxIO Usage (SPARC) - Solaris 11

```
# host_config -setup -protocol fcp -multipath mpzio

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####
```

```
ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:64, throttle-min:8";
```

```
#####
#####
```

```
The following lines will be REMOVED from the /etc/driver/drv/
scsi_vhci.conf file
```

```
#####
#####
```

```
scsi-vhci-failover-override =
"NETAPP LUN", "f_sym";
```

```
Do you want to continue (y/n): y
```

```
#####
To complete the configuration, please run the following commands:
#####
```

```
/usr/sbin/stmsboot -D fp -e      (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

Native Driver with DMP and ALUA Usage (SPARC)

```
# host_config -setup -protocol fcp -multipath dmp
```

```
#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####
```

```
ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:8, throttle-min:2";
```

```
Do you want to continue (y/n): y
```

```
#####
To complete the configuration, please run the following commands:
#####
```

```
/usr/sbin/stmsboot -D fp -d      (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

iSCSI with DMP and No ALUA Usage (SPARC)

```
# host_config -setup -protocol iscsi -multipath dmp -noalua
```

```
#####
```

```

The following lines will be ADDED to the /kernel/drv/sd.conf file
#####
sd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:8, throttle-min:2";

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D iscsi -d (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6

```

iSCSI with MPxIO and ALUA Usage (SPARC) - Solaris 10

```

# host_config -setup -protocol iscsi -multipath mpzio

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####

ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:64, throttle-min:8";

#####
#####
The following lines will be REMOVED from the /kernel/drv/
scsi_vhci.conf file
#####
#####

device-type-scsi-options-list =
"NETAPP LUN", "symmetric-option";
symmetric-option = 0x1000000;

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D iscsi -e (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6

```

iSCSI with MPxIO and NO ALUA (SPARC) - Solaris 10

```
# host_config -setup -protocol iscsi -multipath mpxio -noalua

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####

ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:64, throttle-min:8";

#####
####
The following lines will be ADDED to the /kernel/drv/scsi_vhci.conf
file
#####
####

device-type-scsi-options-list =
"NETAPP LUN", "symmetric-option";
symmetric-option = 0x1000000;

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D iscsi -e (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

iSCSI with MPxIO and ALUA Usage (SPARC) - Solaris 11

```
# host_config -setup -protocol iscsi -multipath mpxio

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####

ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-busy:
30, retries-reset:30, retries-notready:300, retries-timeout:10, throttle-
max:64, throttle-min:8";

#####
###
The following lines will be REMOVED from the /etc/driver/drv/
scsi_vhci.conf file
#####
```

```
###

scsi-vhci-failover-override =
"NETAPP LUN", "f_sym";

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D iscsi -e (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

iSCSI with MPxIO and No ALUA Usage (SPARC) - Solaris 11

```
# host_config -setup -protocol iscsi -multipath mpzio -noalua

#####
The following lines will be ADDED to the /kernel/drv/ssd.conf file
#####

ssd-config-list="NETAPP LUN", "physical-block-size:4096, retries-
busy:30, retries-reset:30, retries-notready:300, retries-timeout:
10, throttle-max:64, throttle-min:8";

#####
####
The following lines will be ADDED to the /etc/driver/drv/
scsi_vhci.conf file
#####
####

scsi-vhci-failover-override =
"NETAPP LUN", "f_sym";

Do you want to continue (y/n): y

#####
To complete the configuration, please run the following commands:
#####

/usr/sbin/stmsboot -D iscsi -e (Do not reboot if prompted)
/usr/sbin/shutdown -y -g0 -i 6
```

(Veritas DMP/FC) Tasks for completing the setup of a Veritas DMP stack

To complete the Host Utilities installation when you're using a Veritas DMP stack, you must configure the system parameters.

The tasks you perform vary slightly depending on your driver.

- Solaris native drivers: You must modify the `/kernel/drv/ssd.conf` file for SPARC and `/kernel/drv/sd.conf` for x86
- iSCSI drivers: You must modify the `/kernel/drv/sd.conf` file for SPARC and x86

There are two ways to modify these files:

- Manually edit the files.
- Use the `host_config` command to modify them. This command is provided as part of the Solaris Host Utilities and automatically sets these files to the correct values.

Note: The `host_config` command does not modify the `/kernel/drv/sd.conf` file unless you are using an x86/x64 processor with MPxIO. For more information, see the information on configuring an MPxIO environment.

For a complete list of the host parameters that the Host Utilities recommend you change and an explanation of why those changes are recommended, see the *Host Settings Affected by the Host Utilities* document.

(Veritas DMP) Before you configure the Host Utilities for Veritas DMP

Before you configure the system parameters for a Veritas DMP environment, you need to create backup files.

- Create your own backup of the files you are modifying:
For systems using Solaris native drivers, make a backup of the `/kernel/drv/ssd.conf` file for SPARC and the `/kernel/drv/sd.conf` for x86.

The `host_config` command automatically creates backups for you, but you can revert to those backups only once. By manually creating the backups, you can revert to them as needed.

Related information

[Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations - www.ibm.com/storage/support/nseries/](http://www.ibm.com/storage/support/nseries/)

(Veritas DMP) `sd.conf` and `ssd.conf` variables for systems using native drivers

If your system uses Solaris native drivers, you need to modify the values in `/kernel/drv/ssd.conf` file for SPARC and in the `/kernel/drv.sd.conf` file for x86.

Note: Versions of the Host Utilities using native drivers always use single-image `cfmode`. If you are using native drivers and not using single-image mode, change your mode.

The required values are:

- `throttle_max=8`
- `not_ready_retries=300`
- `busy_retries=30`
- `reset_retries=30`
- `throttle_min=2`
- `timeout_retries=10`
- `physical_block_size=4096`

The Solaris Host Utilities provides a best-fit setting for target and LUN queue depths.

Tasks for completing the setup of an MPxIO stack

To complete the configuration when you're using an MPxIO stack, you must modify the parameters in either `/kernel/drv/ssd.conf` or `/kernel/drv/sd.conf` and set them to the recommended values.

To set the recommended values, you can either:

- Manually edit the file for your system.
- Use the `host_config` command to automatically make the changes.

For a complete list of the host parameters that the Host Utilities recommend you change and an explanation of why those changes are recommended, see the *Host Settings Affected by the Host Utilities* document.

Before configuring system parameters on an MPxIO stack

Before you configure the system parameters on an MPxIO stack using FC, you need to perform certain tasks.

- Create your own backup of the file you are modifying:
 - `/kernel/drv/ssd.conf` for systems using SPARC processors
 - `/kernel/drv/sd.conf` for systems using x86/x64 processors

The `host_config` command automatically creates backups for you, but you can only revert to those backups once. By manually creating the backups, you can revert to them as many times as needed.

- If MPxIO was previously installed using the Host Utilities or a Host Attach Kit before 3.0.1 and ALUA was not enabled, you must remove it.

Note: iSCSI is not supported with ALUA if you are running Data ONTAP operating in 7-Mode or Data ONTAP before 8.1.1. operating in Cluster-Mode. ALUA is supported in the iSCSI Solaris Host Utilities 3.0 and the Solaris Host Utilities using the FC protocol. However, it is not supported with the iSCSI protocol for the Host Utilities 5.x, the iSCSI Solaris Host Utilities 3.0.1, or Solaris 10 Update 3.

Parameter values for systems using MPxIO

You can manually set the parameter values for systems using MPxIO with the FC protocol by modifying `/kernel/drv/ssd.conf` (SPARC processor systems) or `/kernel/drv/sd.conf` (x86/x64 processor systems).

Both SPARC processor systems and x86/x64 processor systems using MPxIO use the same values. The required values are:

- `throttle_max=64`
- `not_ready_retries=300`
- `busy_retries=30`
- `reset_retries=30`
- `throttle_min=8`
- `timeout_retries=10`
- `physical_block_size=4096`

You must also set the VIP/PID information to “IBM LUN”. You can use the `host_config` command to configure this information.

(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.

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 or later, 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 recommended values for Host Utilities. The section *(Veritas DMP) Setting the restore daemon interval* contains information the values you should use.

(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 Website 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 N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9). This information is updated frequently. After you know the version you need, go to the Symantec Website and download the ASL and 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.

(Veritas DMP) Information provided by the ASL

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

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

- Enclosure name.
With enclosure-based naming, the name of the Veritas disk contains the model name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to VxVM about SAN-attached storage systems, instead of referring to them as Just a Bunch of Disks (JBOD) devices or raw devices. The enclosure-based naming feature used by VxVM creates a disk name based on the name of its enclosure, or disk array, and not a raw device name.
- Multipathing policy. The storage is accessed as either an active/active (A/A-NETAPP) disk array or an active/passive concurrent (A/P-C-NETAPP) disk array. The ASL also provides information about primary and secondary paths to the storage.

For details about system management, see *Veritas Volume Manager Administrator's Guide*. Veritas documents are available at [Veritas Storage Foundation DocCentral](#).

(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 or later, other versions of Veritas Storage Foundation require that you install them.

If you are using Veritas Storage Foundation 5.0 or later, 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 N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9) 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 newer versions 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.

You can use the `pkgrm` command to uninstall the ASL and APM.

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

- Download 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.

(Veritas DMP) ASL and APM installation overview

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

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

- Verify that your configuration meets system requirements. See the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9) for current information about the system requirements.
- If you currently have the ASL installed, determine its version.
- If you need to install a newer version of the ASL and APM, remove the older versions before you install the new versions.

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

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

- Obtain the new ASL and the APM.
- Follow tech note instructions from Symantec (Veritas) to install new versions of ASL and APM. to install the new version of the ASL and APM.

Related information

[N series interoperability matrix website - www.ibm.com/systems/storage/network/interphome.html](http://www.ibm.com/systems/storage/network/interphome.html)

(Veritas) 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
=====
libvxCLARiiON.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
libvxengenio.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
libvxhpalua.so                         vm-5.0-rev-1         5.0
libvxibm4k.so                          vm-5.0-rev-1         5.0
libvxibm6k.so                          vm-5.0-rev-1         5.0
libvxibm8k.so                          vm-5.0-rev-1         5.0
libvxibmsvc.so                         vm-5.0-rev-1         5.0
libvxibm4k.so                          vm-5.0-rev-1         5.0
libvxibm6k.so                          vm-5.0-rev-1         5.0
libvxibm8k.so                          vm-5.0-rev-1         5.0
libvxibmsvc.so                         vm-5.0-rev-1         5.0
libvxnetapp.so                         vm-5.0-rev-0         5.0
```

(Veritas) How to get the ASL and APM

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

To determine which versions of the ASL and APM you need for your version of the host operating system, check the www.ibm.com/systems/storage/network/interophome.html. This information is updated frequently. When you know which version you need, go to the Symantec Website and download the ASL and APM.

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

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

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

Information about getting the Symantec TechNote for the ASL and APM is provided on the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

(Veritas DMP) Installing the ASL and APM software

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

Before you begin

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

Steps

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

```
vxddladm rmjbot vid=NETAPP
```
3. Verify that you have downloaded the correct version of the ASL and APM by checking the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9). If you do not already have the correct version or the ASL and APM TechNote, you can follow the link in the matrix to the correct location on the Symantec Website.

4. Install the ASL and APM according to the installation instructions provided by the ASL/APM TechNote on the Symantec Website.

5. If your host is connected to N series storage, verify your installation by entering:

```
vxdlmpadm listenclosure all
```

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

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

```
vxddladm listsupport all
```

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

```
vxdlmpadm listapm all
```

After you finish

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

- If you have Data ONTAP 7.1 or later, it is recommended that you change the cfmode setting of your clustered systems to single-image mode, and then reconfigure your host to discover the new paths to the disk.

- On the storage system, create LUNs and map them to the igroups containing the WWPNs of the host HBAs.
- On the host, discover the new LUNs and configure them to be managed by VxVM.

Related information

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

(Veritas DMP) Tasks to perform before you uninstall the ASL and APM

Before you uninstall the ASL and APM, you should perform certain tasks.

- Quiesce I/O
- Deport the disk group

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

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

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

```
# pkginfo | grep VRTSNTAP
system          VRTSNTAPapm          Veritas NetApp Array
Policy Module.
system          VRTSNTAPasl          Veritas NetApp Array
Support Library
# pkgrm VRTSNTAPapm

The following package is currently installed:
  VRTSNTAPapm  Veritas NetApp Array Policy Module.
              (sparc) 5.0,REV=09.12.2007.16.16

Do you want to remove this package? [y,n,?,q] y

## Removing installed package instance "VRTSNTAPapm"

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package "VRTSNTAPapm" dependencies in global zone
## Processing package information.
## Executing preremove script.
Check if Module is loaded
## Removing pathnames in class "none"
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm/sparcv9 "shared pathname not removed"
```

```

/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm "shared pathname not removed"
/kernel/drv "shared pathname not removed"
/kernel "shared pathname not removed"
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/64 "shared pathname not removed"
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/32 "shared pathname not removed"
/etc/vx/apmkey.d "shared pathname not removed"
/etc/vx "shared pathname not removed"
/etc "shared pathname not removed"
## Updating system information.

Removal of "VRTSNTAPapm" was successful.
#

# pkgrm VRTSNTAPas1

The following package is currently installed:
  VRTSNTAPas1  Veritas NetApp Array Support Library
               (sparc) 5.0,REV=11.19.2007.14.03

Do you want to remove this package? [y,n,?,q] y

## Removing installed package instance "VRTSNTAPas1"

This package contains scripts which will be executed with super-user
permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?,q] y
## Verifying package "VRTSNTAPas1" dependencies in global zone
## Processing package information.
## Executing preremove script.
Unloading the library
## Removing pathnames in class "none"
/etc/vx/lib/discovery.d/libvxnetapp.so.2
/etc/vx/lib/discovery.d "shared pathname not removed"
/etc/vx/lib "shared pathname not removed"
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/aslkey.d "shared pathname not removed"
## Executing postremove script.
## Updating system information.

```

```
Removal of "VRTSNTAPas1" was successful.
```

(Veritas DMP) Example of installing the ASL and the APM

The following is a sample installation of the ASL and the APM when you have Veritas Storage Foundation 5.0.

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

```
# pkgadd -d . VRTSNTAPas1

Processing package instance "VRTSNTAPas1" from "/tmp"

Veritas NetApp Array Support Library(sparc) 5.0,REV=11.19.2007.14.03
Copyright © 1990-2006 Symantec Corporation. All rights reserved.

Symantec and the Symantec Logo are trademarks or registered
trademarks of
Symantec Corporation or its affiliates in the U.S. and other
countries. Other
names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be
"commercial computer
software" and "commercial computer software documentation" as
defined in FAR
Sections 12.212 and DFARS Section 227.7202.
Using "/etc/vx" as the package base directory.
## Processing package information.
## Processing system information.
   3 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of "VRTSNTAPas1"
[y,n,?] y

Installing Veritas NetApp Array Support Library as "VRTSNTAPas1"

## Installing part 1 of 1.
/etc/vx/aslkey.d/libvxnetapp.key.2
/etc/vx/lib/discovery.d/libvxnetapp.so.2
[ verifying class "none" ]
## Executing postinstall script.
Adding the entry in supported arrays
```

Loading The Library

```

Installation of "VRTSNTAPasl" was successful.
#

# pkgadd -d . VRTSNTAPapm

Processing package instance "VRTSNTAPapm" from "/tmp"
Veritas NetApp Array Policy Module.(sparc) 5.0,REV=09.12.2007.16.16
  Copyright 1996-2005 VERITAS Software Corporation. All rights
  reserved.
  VERITAS, VERITAS SOFTWARE, the VERITAS logo and all other
  VERITAS
  product names and slogans are trademarks or registered
  trademarks of VERITAS Software Corporation in the USA and/or
  other
  countries. Other product names and/or slogans mentioned
  herein may
  be trademarks or registered trademarks of their respective
  companies.

Using "/" as the package base directory.
## Processing package information.
## Processing system information.
   9 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.

This package contains scripts which will be executed with super-user
permission during the process of installing this package.

Do you want to continue with the installation of "VRTSNTAPapm"
[y,n,?] y

Installing Veritas NetApp Array Policy Module. as "VRTSNTAPapm"

## Installing part 1 of 1.
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/32/dmpnetapp.key.SunOS_5.9
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.10
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.8
/etc/vx/apmkey.d/64/dmpnetapp.key.SunOS_5.9
/kernel/drv/vxapm/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/dmpnetapp.SunOS_5.9
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.10
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.8
/kernel/drv/vxapm/sparcv9/dmpnetapp.SunOS_5.9
[ verifying class "none" ]
## Executing postinstall script.

```

```
Installation of "VRTSNTAPapm" was successful.
```

(Veritas DMP) What an ASL array type is

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

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

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

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

(Veritas DMP) The storage system's FC failover mode or iSCSI configuration and the array types

In clustered storage configurations, the array type corresponds to the storage system `cfmode` settings or the iSCSI configuration.

If you use the standby `cfmode` or iSCSI configuration, the array type will be A/A-NETAPP; otherwise, it will be A/P-C-NETAPP.

Note: The ASL also supports direct-attached, non-clustered configurations, including NearStore models. These configurations have no `cfmode` settings. ASL reports these configurations as Active/Active (A/A-NETAPP) array types.

(Veritas DMP) Using VxVM to display available paths

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

Steps

1. View all the devices by entering:

vxdisk list

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

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

```
vxddmpadm getsubpaths dmpnodename=device
```

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

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

```
vxddmpadm getsubpaths ctrl=controller_name
```

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

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

(Veritas) Displaying multipathing information using sanlun

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

About this task

When the ASL is installed and the LUN is controlled by VxVM, the output of the `sanlun` command displays the `Multipath_Policy` as either A/P-C or A/A.

Step

1. On the host, enter the following command:

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

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) Veritas environments and the fast recovery feature

Whether you need to enable or disable the Veritas Storage Foundation 5.0 fast recovery feature depends on your environment.

For example, if your host is using DMP for multipathing and running Veritas Storage Foundation 5.0 with the APM installed, you must have fast recovery **enabled**.

However, if your host is using MPxIO with Veritas, then you must have fast recovery **disabled**.

For details on using fast recovery with different Host Utilities Veritas environments, see the *Solaris Host Utilities 6.1 Release Notes*.

(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 are 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 for 5.0 MP3 and later

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. Change the restore daemon setting to 60 and set the policy to

`check_disabled`

```
/usr/sbin/vxdmpadm settune dmp_restore_interval=60
/usr/sbin/vxdmpadm settune dmp_restore_policy=check_disabled
```

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

2. Verify the changes.

```
/usr/sbin/vxdmpadm gettune dmp_restore_interval
/usr/sbin/vxdmpadm gettune dmp_restore_policy
```

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

```
# vxdmpadm gettune dmp_restore_interval
-----
Tunable                Current Value  Default Value
-----
dmp_restore_interval    60            300

# vxdmpadm gettune dmp_restore_policy
-----
Tunable                Current Value  Default Value
-----
dmp_restore_policy     check_disabled check_disabled
```

(Veritas DMP) Probe Idle LUN settings

Symantec requires that the probe idle lun setting be disabled in versions 5.0 MP3 and later. I/Os are not issued on LUNs affected by controller failover, and during error analysis they are marked as idle. If the probe idle LUN setting is enabled, DMP proactively checks LUNs that are not carrying I/O by sending SCSI inquiry probes. The SCSI inquiry probes performed on paths that are marked idle as a result of controller failover will fail, causing DMP to mark the path as failed.

Steps

1. Execute the following command to disable the setting.

```
/usr/sbin/vxdmpadm settune dmp_probe_idle_lun=off
```

2. Execute the following command to verify the setting.

```
/usr/sbin/vxdmpadm gettune dmp_probe_idle_lun
```

Below is a sample of the output displayed by the above command.

```
# vxdmpadm gettune dmp_probe_idle_lun
-----
Tunable                Current Value  Default Value
-----
dmp_probe_idle_lun
```

(Veritas DMP) DMP Path Age Settings

If the state of the LUN path changes too quickly, DMP will mark the path as suspect. After the path is marked as suspect, it will be monitored and not be used for I/O for the duration of the `dmp_path_age`. The default monitor time is 300 seconds. Starting in 5.1 SP1, Symantec recommends reducing the default time to 120 seconds to allow for quicker recovery.

About this task

Note: These steps apply to 5.1 SP1 and later.

Steps

1. Execute the following command to disable the setting.

```
/usr/sbin/vxdmpadm settune dmp_path_age=120
```

2. Execute the following command to verify the setting.

```
/usr/sbin/vxdmpadm gettune dmp_path_age
```

This is a sample of the output displayed by the above command:

```
# vxdmpadm gettune dmp_path_age
          Tunable
-----
dmp_path_age          120          300
```

(Veritas) Information about ASL error messages

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

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

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

LUN configuration and the Solaris Host Utilities

Configuring and managing LUNs involves several tasks. Whether you are executing the Host Utilities in a Veritas DMP environment or an MPxIO environment determines which tasks you need to perform. The following sections provide information on working with LUNs in all the Host Utilities environments.

Overview of LUN configuration and management

LUN configuration and management involves a number of tasks.

The following table summarizes the tasks for all the supported Solaris environments. If a task does not apply to all environments, the table specifies the environments to which it does apply. You need to perform only the tasks that apply to your environment.

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. After you create the igroup, you must create LUNs on the storage system, and map the LUNs to the igroup.</p> <p>For complete information, refer to your version of the <i>SAN Administration Guide</i> (called <i>Block Access Management Guide for iSCSI and FC</i> in Data ONTAP 8.1 and earlier).</p>
2. (MPxIO) Enable ALUA	<p>If your environment supports ALUA, you must have it set up to work with igroups. To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.</p>
3. (MPxIO, Solaris native drivers with Veritas DMP) Display a list of controllers	<p>If you are using an MPxIO stack or Solaris native drivers with Veritas DMP, you need to get information about the controller before you can discover the LUNs. Use the <code>cfgadm -al</code> command to display a list of controllers.</p>

Task	Discussion
4. Discover LUNs	<p>(iSCSI) When you map new LUNs to the Solaris host, run the following command on the host console to discover the LUNs and create iSCSI device links:</p> <pre>devfsadm -i iscsi</pre> <p>(MPxIO, Solaris native drivers with Veritas) To discover the LUNs, use the command:</p> <pre>/usr/sbin/cfgadm -c configure cx</pre> <p>x is the controller number of the HBA where the LUN is expected to be visible</p>
5. Label LUNs, if appropriate for your system	<p>Use the Solaris format utility to label the LUNs. For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume.</p>
6. Configure volume management software	<p>You must configure the LUNs so they are under the control of a volume manager (SVM, ZFS, or VxVM). Use a volume manager that is supported by your Host Utilities environment.</p>

Related information

The N series support website: www.ibm.com/storage/support/nseries/

Tasks necessary for creating and mapping LUNs

Before you can work with LUNs, you must set them up.

To set LUNs up, do the following:

- Create an igroup.

Note: If you have an active/active configuration, you must create a separate igroup on each system in the configuration.

- Create one or more LUNs and map the LUNs to an igroup.

How the LUN type affects performance

The value you specify for the `ostype` parameter when you create a LUN can affect performance.

For optimal performance, slices or partitions of LUNs must be aligned with the WAFL volume. To achieve optimal performance, you need to provide the correct value for `ostype` for your system.

There are two values for `ostype`:

- `solaris`
- `solaris_efi`

Use the `solaris` `ostype` with UFS and VxFS file systems. When you specify `solaris` as the value for `ostype` parameter, slices or partitions of LUNs are automatically aligned with the WAFL volume.

Solaris uses a newer labeling scheme, known as EFI, for LUNs that will be 2 TB or larger. For these situations, you specify `solaris_efi` as the value for the `ostype` parameter when you create the LUN. If the `solaris_efi` `ostype` is not available, you must perform special steps to align the partitions to the WAFL volume. See the *Solaris Host Utilities Release Notes* for details.

For LUNs used with ZFS zpools, always use the `solaris` `ostype` regardless of the size of the LUN.

Related references

[LUN types, OS label, and OS version combinations for achieving aligned LUNs](#) on page 133

Methods for creating igroups and LUNs

There are several methods for creating igroups and LUNs.

You can create igroups and LUNs on a storage system by entering the following command(s) on the storage system:

- `lun setup`
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
- A series of individual commands such as `lun create`, `igroup create`, and `lun map`
You can use this method to create one or more LUNs and igroups in any order.

Best practices for creating igroups and LUNs

There are several best practices you should consider when you create igroups and LUNs.

The best practices include:

- Disable scheduled snapshots.

- Map the igroup to an application. Make sure the igroup includes all the initiators that the application uses to access its data. (Multiple applications can use the same initiators.)
- Do not put LUNs in the root volume of a storage system. The default root volume is `/vol/vol0`.

(iSCSI) Discovering LUNs

The method you use to discover new LUNs when you are using the iSCSI protocol depends on whether you are using iSCSI with MPxIO or Veritas DMP.

Step

1. To discover new LUNs when you are using the iSCSI protocol, execute the commands that are appropriate for your environment.

- **(MPxIO)** Enter the command:

```
/usr/sbin/devfsadm -i iscsi
```

- **(Veritas)** Enter the commands:

```
/usr/sbin/devfsadm -i iscsi
/usr/sbin/vxdctl enable
```

The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks. You can use the `format` command to label the disk.

Note: Occasionally the `/usr/sbin/devfsadm` command does not find LUNs. If this occurs, reboot the host with the reconfigure option (`touch /reconfigure; /sbin/init 6`)

Solaris native drivers and LUNs

There are several tasks you need to perform when using Solaris native drivers and working with LUNs. The following sections provide information about those tasks.

(Solaris native drivers) Getting the controller number

Before you discover the LUNs, you need to determine what the controller number of the HBA is.

About this task

You must do this regardless of whether you are using Solaris native drivers with MPxIO or Veritas DMP.

Step

1. Use the `cfgadm -al` command to determine what the controller number of the HBA is. If you use the `/usr/sbin/cfgadm -c configure c x` command to discover the LUNs, you need to replace `x` with the HBA controller number.

The following example uses the `cfgadm -al` command to determine the controller number of the HBA. To make the information in the example easier to read, the key lines in the output are shown in **bold**.

```
$ cfgadm -al
Ap_Id                Type                Receptacle
Occupant             Condition
c0                 fc-fabric         connected
configured      unknown
c0::500a098187f93622 disk             connected
configured      unknown
c0::500a098197f93622 disk             connected
configured      unknown
c1                   scsi-bus           connected
configured           unknown
c1::dsk/c1t0d0       disk               connected
configured           unknown
c1::dsk/c1t1d0       disk               connected
configured           unknown
c2                 fc-fabric         connected
configured      unknown
c2::500a098287f93622 disk             connected
configured      unknown
c2::500a098297f93622 disk             connected
configured      unknown
```

(Solaris native drivers) Discovering LUNs

You must both ensure that the host discovers the new LUNs and validate that the LUNs are visible on the host.

About this task

You must do this regardless of whether you are using Solaris native drivers with MPxIO or with Veritas DMP.

Step

1. To discover new LUNs, enter:

```
/usr/sbin/cfgadm -c configure c x
```

where `x` is the controller number of the HBA where the LUN is expected to be visible.

If you do not see the HBA in the output, check your driver installation to make sure it is correct.

The system probes for new devices. When it finds the new LUNs, it might generate a warning about a corrupt label. This warning means that the host discovered new LUNs that need to be labeled as Solaris disks.

Labeling the new LUN on a Solaris host

You can use the `format` utility to format and label new LUNs. This utility is a menu-driven script that is provided on the Solaris host. It works with all the environments supported by the Host Utilities.

Steps

1. On the Solaris host, enter:

```
/usr/sbin/format
```

2. At the `format>` prompt, select the disk you want to modify
3. When the utility prompts you to label the disk, enter `y`. The LUN is now labeled and ready for the volume manager to use.
4. When you finish, you can use the `quit` option to exit the utility.

The following examples show the type of output you would see on a system using LPFC drivers and on a system using Solaris native drivers.

Example 1: This example labels disk number 1 on a system using LPFC drivers. (Portions of this example have been removed to make it easier to review.)

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c3t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
    /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@0,0
  1. c4t0d0 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
    /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,0
  2. c4t0d1 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
    /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,1
  3. c4t0d2 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
    /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,2
  4. c4t0d3 <NETAPP-LUN-7310 cyl 1232 alt 2 hd 16 sec 128>
    /pci@7c0/pci@0/pci@1/pci@0,2/lpfc@1/sd@0,3

Specify disk (enter its number): 1
selecting c4t0d0
[disk formatted]
```

```
...
Disk not labeled. Label it now? y
```

Example 2: This example labels disk number 2 on a system that uses Solaris native drivers with Veritas DMP. (Portions of this example have been removed to make it easier to review.)

```
$ format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e01008eb71,0
  1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
    /pci@8,600000/SUNW,qlc@4/fp@0,0/ssd@w500000e0100c6631,0
  2. c6t500A098387193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd
16 sec 2048>
    /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098387193622,0
  3. c6t500A098197193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd
16 sec 2048>
    /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098197193622,0
  4. c6t500A098187193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd
16 sec 2048>
    /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098187193622,0
  5. c6t500A098397193622d0 <NETAPP-LUN-0.2 cyl 6398 alt 2 hd
16 sec 2048>
    /pci@8,600000/emlx@1/fp@0,0/ssd@w500a098397193622,0

Specify disk (enter its number): 2
selecting c6t500A098387193622d0: TESTER
[disk formatted]
...
Disk not labeled. Label it now? y
```

Example 3: This example runs the `fdisk` command and then labels disk number 15 on an x86/x64 system. You must run the `fdisk` command before you can label a LUN.

```
$Specify disk (enter its number): 15
selecting C4t60A9800043346859444A2D367047492Fd0
[disk formatted]

FORMAT MENU:
disk          - select a disk
type         - select (define) a disk type
partition    - select (define) a partition table
current      - describe the current disk
format       - format and analyze the disk
fdisk        - run the fdisk program
repair       - repair a defective sector
label        - write label to the disk
analyze      - surface analysis
defect       - defect list management
backup       - search for backup labels
verify       - read and display labels
save         - save new disk/partition definitions
```

```

    inquiry    - show vendor, product and revision
    volname    - set 8-character volume name
    !>cmd>    - execute >cmd>, then return
    quit
format> label
Please run fdisk first.
format> fdisk
No fdisk table exists. The default partition for the disk is:

    a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to
edit the
partition table.
y
format> label
Ready to label disk, continue? y

format>

```

Methods for configuring volume management

When your configuration uses volume management software, you must configure the LUNs so they are under the control of the volume manager.

The tools you use to manage your volumes depend on the environment you are working in: Veritas DMP or MPxIO.

Veritas DMP If you are a Veritas DMP environment, even if you are using Solaris native drivers or the iSCSI protocol, you must use VxVM to manage the LUNs. You can use the following Veritas commands to work with LUNs:

- The Veritas `/usr/sbin/vxdctl enable` command brings new LUNs under Veritas control.
- The Veritas `/usr/sbin/vxdiskadm` utility manages existing disk groups.

MPxIO If you are in an MPxIO environment, you can manage LUNs using SVM, ZFS, or, in some cases, VxVM.

Note: To use VxVM in an MPxIO environment, first check the N series support website (accessed and navigated as described in [Websites](#) on page 9) to see if your environment supports VxVM.

For additional information, refer to the documentation that shipped with your volume management software.

Related information

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

The sanlun utility

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

Displaying host LUN information with `sanlun`

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

Steps

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

```
cd /opt/ontap/SANToolkit/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 | controller/vserver_name | controller/vserver_name:<path_name>]
```

`-v` produces verbose output.

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

- *host device filename* is the special device file name for the disk on Solaris (which might represent a storage system LUN).
- `all` lists all storage system LUNs under `/dev/rdisk`.
- *controller/vserver_name* lists all storage system LUNs under `/dev/rdisk` on that storage system.
- *controller/vserver_name:path_name* lists all storage system LUNs under `/dev/rdisk` that are connected to the storage system path name LUN on that storage system.

`-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 -v [ all|controller/vserver_name| controller/vserver_name:path_Name ]
```

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.

The requested LUN information is displayed.

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/rdisk/<x>`
to display the summary listing of the LUN(s) associated with the host device `/dev/rdisk<x>` where `x` is a device such as `/dev/rdisk/c##d#`.
- `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: When you specify either the `sanlun lun show <storage_system_name>` or the `sanlun lun show <storage_system_name:storage_system_pathname>` command, the utility displays only the LUNs that have been discovered by the host. LUNs that have not been discovered by the host are not displayed.

The following is an example of the output you see when you use the `sanlun show` command in verbose mode on a system using the Veritas DMP stack.

```
# ./sanlun lun show -v
filer: lun-pathname          device filename      adapter lun
size      lun state
filerX:  /vol/vol1/hostA_lun2 /dev/rdisk/c0t500A098487093F9Dd5s2
emlxs0   3g (3221225472) GOOD
          Serial number: HnTMWZEgHyD5
          Filer FCP nodename:500a098087093f9d Filer FCP portname:
00a098397093f9d
          Filer adapter name: 0a
          Filer IP address:      10.60.181.66
          Filer volume name:vol1  FSID:0x331ee81
          Filer qtree name:/vol/vol1  ID:0x0
          Filer snapshot name:    ID:0x0
          LUN partition table permits multiprotocol access: no
                                why: bad starting cylinder or bad size for data
```

```
partition
  LUN has valid label: yes
```

Displaying path information with sanlun

You can use `sanlun` to display information about the paths to the storage system.

Steps

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

```
sanlun lun show -p
```

`-p` provides information about the optimized (primary) and non-optimized (secondary) paths available to the LUN when you are using multipathing.

Note: (MPxIO stack) MPxIO makes the underlying paths transparent to the user. It only exposes a consolidated device such as `/dev/rdsk/c7t60A980004334686568343655496C7931d0s2`. This is the name generated using the LUN's serial number in the IEEE registered extended format, type 6. The Solaris host receives this information from the SCSI Inquiry response. As a result, `sanlun` cannot display the underlying multiple paths. Instead it displays the target port group information. You can use the `mpathadm` or `luxadm` command to display the information if you need it.

`all` lists all storage system LUNs under `/dev/rdsk`.

Explanation of the sanlun lun show -p output

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

- **(Veritas DMP stack) path state**—Whether the path is enabled or disabled.
- **(Veritas DMP stack) path type:**
 - 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. Note that this case occurs only when Veritas assumes the array policy is Active/Active. If the array policy is Active/Passive and the path is being served by the partner file in the takeover mode, that path state displays as secondary.
- **device filename**—The special device file name for the disk on Solaris that represents the LUN.
- **host HBA**—The name of the initiator HBA on the host.

- local storage system port—The port that provides direct access to a LUN. This port appears as the on N series storage systems, where X is the slot number of the HBA. On N5200 and N5500 storage systems, the port is Xa. This is always a primary (optimized) path.
- partner storage system port—The port that provides passive path failover. This port appears as X_C on N series storage systems, where X is the slot number of the HBA. This is always a secondary (unoptimized or indirect) 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.

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 instructions.

-v option produces verbose output.

all lists information for all FC adapters.

The FC adapter information is displayed.

The following command line displays information about the adapter on a system using the qlc driver.

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

adapter name: qlc0
WWPN: 21000003ba16dec7
WWNN: 20000003ba16dec7
driver name: 20060630-2.16
model: 2200
model description: 2200
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 2.1.144
Number of ports: 1
port type: Private Loop
port state: Operational
supported speed: 1 GBit/sec
negotiated speed: 1 GBit/sec
```

```
OS device name: /dev/cfg/c1
adapter name: qlc1
WWPN: 210000e08b88b838
WWNN: 200000e08b88b838
driver name: 20060630-2.16
model: QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 4.0.23
Number of ports: 1 of 2
port type: Fabric
port state: Operational
supported speed: 1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed: 4 GBit/sec
OS device name: /dev/cfg/c2
adapter name: qlc2
WWPN: 210100e08ba8b838
WWNN: 200100e08ba8b838
driver name: 20060630-2.16
model: QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number: Not Available
hardware version: Not Available
driver version: 20060630-2.16
firmware version: 4.0.23
Number of ports: 2 of 2
port type: Fabric
port state: Operational
supported speed: 1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed: 4 GBit/sec
OS device name: /dev/cfg/c3
```

About the collectinfo command

`collectinfo` is an optional command that collects anonymous support information about system hosts such as HBA types, volume manager configuration and operating system.

When `collectinfo` is executed, this information is pushed to the IBM N series controller. If the IBM N series controller has AutoSupport enabled, this information is part of the payload that is used to ensure that future releases of the Host Utilities meet customer needs. This command must be initiated by the user.

Note: You must be logged as "root" to run `collectinfo`.

The `collectinfo` command has the following format:

- `collectinfo [-controller 7-Mode_Controller_IP_or_Hostname] [-user username] [-password password] [-use_http]`
- `collectinfo [-controller Cluster-Mode_Data_Vserver_IP_or_Hostname] [-user username] [-password password] [-use_http]`
- `collectinfo [-help|-h|-?]`
- `collectinfo [-version]`

collectinfo options

The `collectinfo` command has several options you can use. These options apply to all environments. This command is executed on the host.

Option	Description
<code>-controller Cluster-Mode_Data_Vserver_IP_or_Hostname 7-Mode_Controller_IP_or_Hostname></code>	<p>For Data ONTAP operating in Cluster-Mode: Specify the IP address or hostname of a LIF on the cluster Vserver. This LIF should have data-protocol set to "none" and cannot be a cluster-mgt LIF.</p> <p>For Data ONTAP operating in 7-Mode: Specify the IP address or hostname of one of the controllers. You will be prompted for this information if it is not specified.</p>
<code>-user <username></code>	Username to connect to the controller. You will be prompted for this information if it is not specified.

Option	Description
<code>-password <password></code>	Password to use when connecting. You will be prompted for this information if it is not specified.
<code>-use_http</code>	Use the HTTP protocol to connect instead of the default protocol of HTTPS. Note: If this option is chosen, information will be sent over the wire to the controller in an unencrypted format.
<code>-help -h -?</code>	Print the usage statement.
<code>-version</code>	Print the program version.

collectinfo command example

The following example takes you through the process of using the `collectinfo` command.

```
# ./collectinfo.pl -controller test01 -user root -password ""
Please enter the password for user 'root':
Gathering system information

General Info:
=====
Hostid:                               162f26f7
OS:                                    Solaris 10U10
Server Platform:                       i86pc
File Systems Mounted:                   UFS ZFS

Multipathing Info:
=====
Native Multipathing:                   In Use
Veritas DMP:                            Not In Use

Native VM and Cluster Info:
=====
Solaris Volume Manager:                 Not In Use
Sun Cluster:                            Not Installed
ZFS File System:                         In Use

Symantec Product Info:
=====
VCS:                                    Not Installed
Veritas Volume Manager:                 Not Installed
Veritas File System:                    Not Installed
Veritas Cluster File System:            Not Installed
```

```
HBA Info:
```

```
=====
```

Model	Driver	Firmware
111-00778	2.50o	2.701.462.1
LPe11002-M4	2.50o	2.70a5

```
Sending config to the EMS logs on controller 'test01'
```

SAN boot LUNs in a Solaris Native FC environment

You can set up a SAN boot LUN to work in a Veritas DMP environment or a Solaris MPxIO environment using the FC protocol and running the Solaris Host Utilities. The method you use to set up a SAN boot LUN can vary depending on your volume manager and file system.

To verify that SAN booting is supported in your configuration, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

If you are using Solaris native drivers, refer to Solaris documentation for details about additional configuration methods. In particular, see the Oracle document, *Sun StorEdge SAN Foundation Software 4.4 Configuration Guide*.

Related information

[Sun StorEdge SAN Foundation Software 4.4 Configuration Guide](#)

Prerequisites for creating a SAN boot LUN

You need to have your system set up and the Host Utilities installed before you create a SAN boot LUN for a Veritas DMP environment.

Note: SAN booting is not supported with the iSCSI protocol.

Before attempting to create a SAN boot LUN, make sure the following prerequisites are in place:

- The Solaris Host Utilities software and supported firmware is installed.
- The host operating system is installed on a local disk and uses a UFS file system.
- Boot code/FCode is downloaded and installed on the HBA.
 - For Emulex HBAs, the FCode is available on the Emulex site.
 - For QLogic-branded HBAs, the FCode is available on the QLogic site.
 - For Oracle-branded QLogic HBAs, the FCode is available as a patch from Oracle.
 - If you are using Emulex HBAs, you must have the Emulex FCA utilities with `emlxdrv`, `emlxadm`, and `hbacmd` commands.
 - If you are using Emulex-branded HBAs or Oracle-branded Emulex HBAs, make sure you have the current FCode, available on the Emulex site.
 - If you are using QLogic-branded HBAs, you must have the SANsurfer SCLI utility installed. You can download the SANsurfer SCLI utility from the QLogic website.

General SAN Boot Configuration Steps

To configure a bootable LUN, you must perform several tasks.

Steps

1. Make sure the HBA is set to the appropriate boot code.
2. Create the boot LUN
 - a) Create the LUN that you will use for the bootable LUN.
 - b) Display the size and layout of the partitions of the current Solaris boot drive.
 - c) Partition the bootable LUN to match the host boot drive.
3. Select the method for installing to the SAN booted LUN.
 - a) If you are using UFS, perform a file system dump and restore to the LUN chosen to be used for SAN boot.
 - b) If you are using ZFS, create a new boot environment using the new LUN.
 - c) Directly install the boot blocks or GRUB information onto the bootable LUN.
4. Modify the boot code.
 - a) Verify the boot code version.
 - b) Set the FC topology to the bootable LUN.
 - c) Bind the adapter target and the bootable LUN for x86 hosts.
 - d) For SPARC hosts, create an alias for the bootable LUN.
5. Reboot the system.

About SPARC OpenBoot

When you are setting up a SAN boot LUN, you can modify OpenBoot to create an alias for the bootable LUN. The alias substitutes for the device address during subsequent boot operations.

OpenBoot is firmware that the host uses to start the system. OpenBoot firmware also includes the hardware-level user interface that you use to configure the bootable LUN.

The steps you need to perform to modify OpenBoot differ depending on whether you are using Solaris native drivers.

About setting up the Oracle native HBA for SAN booting

Part of configuring a bootable LUN when using Veritas DMP with Solaris native drivers is setting up your HBA. To do this, you might need to shut down the system and switch the HBA mode.

SAN booting supports two kinds of Oracle native HBAs. The actions you take to set up your HBA depend on the type of HBA you have.

- If you have an Emulex HBA on a SPARC system, you must make sure the HBA is in SFS mode.
- If you have a QLogic HBA on a SPARC system, you must change the HBA to enable FCode compatibility.

SPARC: Changing the Emulex HBA to SFS mode

To change the mode on an Emulex HBA from an SD compatible mode to an SFS mode, you must bring the system down and then change each HBA.

About this task

Caution: These steps will change the device definition from `lpfc@` to `emlxs@`. Doing this will cause the controller instance to be incremented. Any currently existing devices that are being modified will receive new controller numbers. If you are currently mounting these devices by using the `/etc/vfstab` file, those entries will become invalid.

Steps

1. At the operating system prompt, issue the `init 0` command.

```
# init 0
```

2. When the `ok` prompt appears, enter the `setenv auto-boot? false` command.

```
ok > setenv auto-boot? false
```

3. Enter the `reset-all` command.

```
ok reset-all
```

4. Issue the `show-devs` command to see the current device names.

The following example uses the `show-devs` command to see if the Emulex device has been set to SFS mode. In this case, executing the command shows that the device has not been set to SFS mode because the devices (shown in **bold**) are displayed as `.../lpfc`, not `.../emlxs`. See Step 5 for information setting the devices to SFS mode.

```

ok show-devs
controller@1,400000
/SUNW,UltraSPARC-III+@1,0
/memory-controller@0,400000
/SUNW,UltraSPARC-III+@0,0
/virtual-memory
/memory@m0,0
/aliases
/options
/openprom
/chosen
/packages
/pci@8,600000/SUNW,qlc@4
/pci@8,600000/SUNW,qlc@4/fp@0,0
/pci@8,600000/SUNW,qlc@4/fp@0,0/disk
/pci@8,700000/lpfc@3
/pci@8,700000/lpfc@1
/pci@8,700000/scsi@6,1
/pci@8,700000/scsi@6
/pci@8,700000/usb@5,3

```

5. Select the first Emulex device and set it to SFS mode using the `set-sfs-boot` command. Doing this changes the devices to emlxs devices.

In this example, the `select` command selects the device `lpfc@0`. The `set-sfs-boot` command sets the HBA to SFS mode.

```

ok select /pci@8,700000/lpfc@1

ok set-sfs-boot
Flash data structure updated.
Signature 4e45504f
Valid_flag 4a
Host_did 0
Enable_flag 5
SFS_Support 1
Topology_flag 0
Link_Speed_flag 0
Diag_Switch 0
Boot_id 0
Lnk_timer f
Plogi-timer 0
LUN (1 byte) 0
DID 0
WWPN
LUN (8 bytes)
0000.0000.0000.0000
0000.0000.0000.0000

*** Type reset-all to update. ***
ok

```

6. Repeat Step 5 for each Emulex device.
7. Enter the `reset-all` command to update the devices.

In this example, the `reset-all` command updates the Emulex devices with the new mode.

```
ok reset-all
Resetting ...
```

8. Issue the `show-devs` command to confirm that you have changed the mode on all the Emulex devices.

The following example uses the `show-devs` command to confirm that the Emulex devices are showing up as `emlx` devices. To continue the example shown in from Step 5, the device selected there, `/pci@0/pci@0/pci@8/pci@0/pci@1/lpfc@0`, has been changed to an `emlxs` device. In a production environment, you would want to change all the devices to `emlxs`.

```
ok> show-devs
controller@1,400000
/SUNW,UltraSPARC-III+@1,0
/memory-controller@0,400000
/SUNW,UltraSPARC-III+@0,0
/virtual-memory
/memory@m0,0
/aliases
/options
/openprom
/chosen
/packages
/pci@8,600000/SUNW,qlc@4
/pci@8,600000/SUNW,qlc@4/fp@0,0
/pci@8,600000/SUNW,qlc@4/fp@0,0/disk
/pci@8,700000/emlx@3
/pci@8,700000/emlx@1
/pci@8,700000/scsi@6,1
/pci@8,700000/scsi@6
/pci@8,700000/usb@5,3
```

9. Set the `auto-boot?` back to `true` and boot the system with a reconfiguration boot.

This example uses the `boot` command to bring the system back up.

```
ok setenv auto-boot? true
ok boot -r
```

SPARC: Changing the QLogic HBA to enable FCode compatibility

To enable FCode compatibility on a QLogic HBA, you must bring the system down and then change each HBA.

Steps

1. At the operating system prompt, issue the `init 0` command.

```
# init 0
```

2. When the `ok` prompt is displayed, enter the `setenv auto-boot? false` command.

```
ok setenv auto-boot? false
```

3. Enter the `reset-all` command.

```
ok reset-all
```

4. Issue the `show-devs` command to see the current device names.

The following example uses the `show-devs` command to see whether there is FCode compatibility. The example has been truncated to make it easier to read.

```
ok show-devs
...
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1
/pci@7c0/pci@0/pci@8/QLGC,qlc@0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0,1/fp@0,0/disk
/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0
/pci@7c0/pci@0/pci@8/QLGC,qlc@0/fp@0,0/disk
```

5. Select the first QLogic device.

This example uses the `select` command to select the first QLogic device.

```
ok select /pci@7c0/pci@0/pci@8/QLGC,qlc@0,1
QLogic QLE2462 Host Adapter Driver(SPARC): 1.16 03/10/06
```

6. If you need to set the compatibility mode to FCode, execute the `set-mode` command.

This command is not available nor required for all Qlogic HBAs.

The following example uses the `set-mode` command to set the compatibility mode to FCode.

```
ok set-mode
Current Compatibility Mode: fcode
Do you want to change it? (y/n)
Choose Compatibility Mode:
0 - fcode
1 - bios
enter: 0
Current Compatibility Mode: fcode
```

7. Execute the `set-fc-mode` command to set the FCode mode to `qlc`.

The following example uses the `set-mode` command to set the mode to `qlc`.

```
ok set-fc-mode
Current Fcode Mode: qlc
Do you want to change it? (y/n)
Choose Fcode Mode:
0 - qlc
1 - qla
enter: 0
Current Fcode Mode: glc
```

8. Repeat the previous steps to configure each QLogic device.
9. Enter the `reset-all` command to update the devices.

In this example, the `reset-all` command updates the QLogic devices with the new mode.

```
ok reset-all
Resetting ...
```

10. Set the `auto-boot?` back to true and boot the system with a reconfiguration boot.

This example uses the `boot` command to bring the system back up.

```
ok setenv auto-boot? true
ok boot -r
```

Information on creating the bootable LUN

After setting up the HBAs, you must create the LUN you want to use as a bootable LUN.

Use standard storage system commands and procedures to create the LUN and map it to a host.

In addition, you must partition the bootable LUN so that it matches the partitions on the host boot device. Partitioning the LUN involves:

- Displaying information about the host boot device.
- Modifying the bootable LUN to model the partition layout of the host boot device.

Veritas DMP Systems: Gathering SAN boot LUN information

Before copying any data, it is important to gather information about the LUN you are going to use for SAN booting. You will need this information to complete the boot process.

Steps

1. Run `sanlun lun show` to get a list of available SAN attached devices.

Example

```
# sanlun lun show
controller(7mode) /

vserver(Cmode)      lun-
pathname
-----
n5500-shu05         /vol/vol213/lun213

device
filename
-----
/dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s2

host      lun
adapter  protocol  size  mode
-----
emlxs2   FCP        80.0g  7
```

2. Run `vxdisk -e list` to get a list of LUNs available to Veritas.

Example

```
# vxdisk -e list
DEVICE      TYPE          DISK          GROUP
disk_0      auto:SVM      -             -
disk_1      auto:SVM      -             -
n55000_0-shu05  auto          -             -
n55000_1-shu05  auto          -             -
n55000_2-shu05  auto          -             -
n55000_3-shu05  auto          -             -
n55000_4-shu05  auto          -             -
n55000_5-shu05  auto          -             -

STATUS      OS_NATIVE_NAME  ATTR
SVM         c1t1d0s2        -
SVM         c1t0d0s2        -
nolabel     c3t500A0986974988C3d239s2  tprclm
nolabel     c3t500A0986874988C3d235s2  tprclm
nolabel     c3t500A0986974988C3d237s2  tprclm
nolabel     c3t500A0986974988C3d241s2  tprclm
nolabel     c3t500A0986974988C3d243s2  tprclm
nolabel     c3t500A0986874988C3d236s2  tprclm
```

3. Choose a LUN and then run `vxdisk list <device>` on the device to get the list of primary and secondary paths.

Example

This example uses LUN `n55000_1`.

```
# vxdisk list fas30700_1
Device:      n55000_1
devicetag:  n55000_1
type:       auto
flags:      nolabel private autoconfig
pubpaths:   block=/dev/vx/dmp/n55000_1 char=/dev/vx/rdmp/n55000_1
guid:       -
udid:       NETAPP%5FLUN%5F30019945%5F1ka8k%5DBZ8yq6
site:       -
errno:      Disk is not usable
Multipathing information:
numpaths:   4
c3t500A0986874988C3d235s2      state=enabled   type=secondary
c3t500A0986974988C3d235s2      state=enabled   type=primary
c2t500A0985874988C3d235s2      state=enabled   type=secondary
c2t500A0985974988C3d235s2      state=enabled   type=primary
```

4. Run the `luxadm display` command to find the device path that is mapped to the WWPN of one of the primary paths.

`luxadm` will not designate primary and secondary status so you will need to use the `vxdisk list` and the `luxadm display` output to determine which paths are primary and which are secondary. The path in this example is in bold below.

Example

```
# luxadm display /dev/rdisk/c3t500A0986974988C3d235s2
DEVICE PROPERTIES for disk: /dev/rdisk/c3t500A0986974988C3d235s2
Vendor:      NETAPP
Product ID:  LUN
Revision:    811a
Serial Num:  1ka8k]BZ8yq6
Unformatted capacity: 5120.000 MBytes
Read Cache:  Enabled
  Minimum prefetch: 0x0
  Maximum prefetch: 0x0
Device Type: Disk device
Path(s):

  /dev/rdisk/c3t500A0986974988C3d235s2
  /devices/pci@7c0/pci@0/pci@9/QLGC,qlc@0,1/fp@0,0/
  ssd@w500a0986974988c3,eb:c,raw
    LUN path port WWN:      500a0986974988c3
    Host controller port WWN: 210100e08ba8bf2b
    Path status:             O.K.
  /dev/rdisk/c2t500A0985874988C3d235s2
  /devices/pci@7c0/pci@0/pci@9/QLGC,qlc@0/fp@0,0/
  ssd@w500a0985874988c3,eb:c,raw
    LUN path port WWN:      500a0985874988c3
    Host controller port WWN: 210000e08b88bf2b
    Path status:             O.K.
  /dev/rdisk/c2t500A0985974988C3d235s2
  /devices/pci@7c0/pci@0/pci@9/QLGC,qlc@0/fp@0,0/
  ssd@w500a0985974988c3,eb:c,raw
```

```

LUN path port WWN:          500a0985974988c3
Host controller port WWN:   210000e08b88bf2b
Path status:                O.K.
/dev/rdsk/c3t500A0986874988C3d235s2
/devices/pci@7c0/pci@0/pci@9/QLGC,qlc@0,1/fp@0,0/
ssd@w500a0986874988c3,eb:c,raw
LUN path port WWN:          500a0986874988c3
Host controller port WWN:   210100e08ba8bf2b
Path status:                O.K.

```

5. Document the device path, WWPN and LUN ID in hex.

Native MPxIO Systems: Gathering SAN boot LUN information

Before copying any data, it is important to gather information about the LUN you are going to use for SAN booting. You will need this information to complete the boot process.

Steps

1. Run `sanlun lun show` to get a list of available SAN-attached devices.

Example

```

# sanlun lun show
controller (7mode) /
device
vserver (Cmode)      lun-pathname
filename
-----
n5500-shu05          /vol/vol1213/lun213 /dev/rdsk/
c6t60A98000316B61386B5D425A38797065d0s2

host                lun
adapter            protocol  size    mode
-----
emlxs2             FCP      80.0g   7

```

2. Run `luxadm display <device>`.
The value of `<device>` should be the `/dev/rdsk` path of the SAN boot LUN.
3. Identify and document the WWPN of a primary path to the LUN, the device path of the HBA, and the LUN ID in hex.

Example

In this example:

- The WWPN = 500a0983974988c3
- The HBA = /dev/cfg/c4
- The LUN ID = d5

These values are highlighted in **bold** below.

```
# luxadm display /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s2
DEVICE PROPERTIES for disk: /dev/rdisk/
c6t60A98000316B61386B5D425A38797065d0s2
Vendor:                NETAPP
Product ID:            LUN
Revision:              811a
Serial Num:            1ka8k]BZ8type
Unformatted capacity: 81926.000 MBytes
Read Cache:           Enabled
  Minimum prefetch:    0x0
  Maximum prefetch:    0x0
Device Type:           Disk device
Path(s):

/dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s2
/devices/scsi_vhci/disk@g60a98000316b61386b5d425a38797065:c,raw
Controller              /dev/cfg/c4
  Device Address          500a0983874988c3,d5
  Host controller port WWN 10000000c96e9854
  Class                   secondary
  State                   ONLINE
Controller              /dev/cfg/c4      Device Address
500a0983974988c3,d5
  Host controller port WWN 10000000c96e9854
  Class                   primary
  State                   ONLINE
Controller              /dev/cfg/c5
  Device Address          500a0984874988c3,d5
  Host controller port WWN 10000000c96e9855
  Class                   secondary
  State                   ONLINE
Controller              /dev/cfg/c5
  Device Address          500a0984974988c3,d5
  Host controller port WWN 10000000c96e9855
  Class                   primary
  State                   ONLINE
```

4. If the HBA device path is dev/cfg/c*, then you need to decode that to the pci device path using the `ls -l <path>` command.

Example

```
# ls -l /dev/cfg/c4
lrwxrwxrwx  1 root      root          61 Nov  6 15:36 /dev/cfg/c4 ->
../../devices/pci@0,0/pci8086,25f9@6/pci10df,fe00@0/fp@0,0:fc
```

In this example, the path you want to use is `/pci@0,0/pci8086,25f9@6/pci10df,fe00@0/fp@0,0`.

Gathering source disk information

Before copying any data, it is important to gather information about the LUN you are going to use for SAN booting. You will need this information to complete the boot process.

Steps

1. Run `df /` for UFS systems.

Example

```
# df /
Filesystem          kbytes    used    avail capacity  Mounted on
/dev/dsk/c1t0d0s0   52499703 7298598 44676108    15%      /
```

2. Run `zpool status rpool` for ZFS systems.

Example

```
# zpool status rpool
pool: rpool
state: ONLINE
scan: none requested
config:

   NAME                STATE          READ WRITE CKSUM
   rpool                ONLINE         0     0     0
     c4t0d0s0           ONLINE         0     0     0

errors: No known data errors
```

3. Run `zfs list -t filesystem -r rpool` for ZFS systems.

Example

```
# zfs list -t filesystem -r rpool
NAME                USED  AVAIL  REFER  MOUNTPOINT
rpool                16.0G  212G   50K    /rpool
```

rpool/ROOT	9.83G	212G	31K	legacy
rpool/ROOT/solaris	6.27M	212G	4.23G	/
rpool/ROOT/solaris/var	3.64M	212G	331M	/var

4. Run `prtvtoc` on the `/dev/rdisk` path for the internal disk.

Example

```
# prtvtoc /dev/rdisk/c1t0d0s2
* /dev/rdisk/c1t0d0s2 partition map
*
* Dimensions:
*   512 bytes/sector
*   424 sectors/track
*   24 tracks/cylinder
*  10176 sectors/cylinder
*  14089 cylinders
*  14087 accessible cylinders
*
* Flags:
*   1: unmountable
*  10: read-only
*
*
* Partition  Tag  Flags      First      Sector      Last
*           Tag  Flags      Sector      Count      Sector      Mount Directory
*   0         2    00      4202688  106613952  110816639  /
*   1         3    01           0    4202688    4202687
*   2         5    00           0  143349312  143349311
*   3         0    00    110816640  29367936  140184575  /altroot
*   6         0    00    140184576   2106432  142291007  /globaldevices
*   7         0    00    142291008   1058304  143349311
```

Partitioning and labeling SAN boot LUNs

After you have selected the LUN that you are going to use for SAN booting, you will need to partition and label the device.

About this task

You should set up the slices on your SAN boot LUN so it matches your source boot LUN. A slice will need to be created on the target LUN for each slice you need to copy from the source LUN. If you have slices on the source LUN that you do not want to copy, then you do not need to create slices for them on the target device.

Note: If you have an X64 host, you need to put an `fdisk` partition on the LUN before you can partition and label it.

Steps

1. Run the `format` command.

Example

```
# format
```

2. Choose your SAN boot LUN from the menu.

Example

```
# format
Searching for disks...done

c6t60A98000316B61386B5D425A38797146d0: configured with capacity of
79.98GB

AVAILABLE DISK SELECTIONS:
  0. c0t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
    /pci@780/pci@0/pci@9/scsi@0/sd@0,0
  1. c0t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
    /pci@780/pci@0/pci@9/scsi@0/sd@1,0
  2. c6t60A98000316B61386B5D425A38797146d0 <NETAPP-LUN-811a cyl
6300 alt 2 hd 16 sec 1664>
    /scsi_vhci/ssd@g60a98000316b61386b5d425a38797146
Specify disk (enter its number): 2
```

3. Choose partition from the next menu.

Example

```
FORMAT MENU:
  disk           - select a disk
  type           - select (define) a disk type
  partition      - select (define) a partition table
  current        - describe the current disk
  format         - format and analyze the disk
  repair         - repair a defective sector
  label          - write label to the disk
  analyze        - surface analysis
  defect         - defect list management
  backup         - search for backup labels
  verify         - read and display labels
  save           - save new disk/partition definitions
  inquiry        - show vendor, product and revision
  volname        - set 8-character volume name
  !<cmd>         - execute <cmd>, then return
format> partition
```

4. Partition the LUN by using the `modify` option and a free-hog slice.

The free hog-slice will contain all remaining available space after all other slices have been configured.

Note: The default free-hog slice is slice 6. If you are using this method for booting, then you should switch to slice 0.

Example

```
partition> modify
Select partitioning base:
    0. Current partition table (default)
    1. All Free Hog
Choose base (enter number) [0]? 1

Part      Tag      Flag      Cylinders      Size      Blocks
  0      root      wm         0              0
(0/0/0)      0
  1      swap      wu         0              0
(0/0/0)      0
  2      backup    wu        0 - 6299      79.98GB    (6300/0/0)
167731200
  3 unassigned    wm         0              0
(0/0/0)      0
  4 unassigned    wm         0              0
(0/0/0)      0
  5 unassigned    wm         0              0
(0/0/0)      0
  6      usr      wm         0              0
(0/0/0)      0
  7 unassigned    wm         0              0
(0/0/0)      0

Do you wish to continue creating a new partition
table based on above table[yes]? yes
Free Hog partition[6]? 0
Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]: 4g
Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]: 2g
Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]: 512m

Part      Tag      Flag      Cylinders      Size      Blocks
  0      root      wm        0 - 5785      73.46GB    (5786/0/0)
154046464
  1      swap      wu       5786 - 6101      4.01GB    (316/0/0)
8413184
  2      backup    wu        0 - 6299      79.98GB    (6300/0/0)
167731200
  3 unassigned    wm         0              0
(0/0/0)      0
  4 unassigned    wm         0              0
(0/0/0)      0
  5 unassigned    wm         0              0
(0/0/0)      0
```

```

6          usr      wm      6102 - 6259          2.01GB      (158/0/0)
4206592
7 unassigned      wm      6260 - 6299          520.00MB    (40/0/0)
1064960

Okay to make this the current partition table[yes]? yes
Enter table name (remember quotes): "SANBOOT"

Ready to label disk, continue? y

```

5. Fill out the requested information for each slice you want to create.

Remember the following:

- You want to create slices for each slice represented on the source device.
- If you are using UFS, you will likely want a slice for swap.
- If you are using ZFS, there is no need to create a slice for swap. Swap will be cut out of the zpool.

Related tasks

[Labeling the new LUN on a Solaris host](#) on page 70

UFS File systems: Copying data from locally booted disk

The `ufsdump` command is used to copy the data from the source disk to the SAN boot LUN for UFS file systems. This could be either a native MPxIO system or a Veritas DMP system.

Steps

1. Use the `newfs` command to place a UFS file system on each slice you want to copy.

Example

This example uses only slice 0.

```
# newfs /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0
```

2. Mount the file system that you will use as your target boot device.

Example

```
# mount /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0 /mnt/bootlun
```

3. Use the `ufsdump` command to copy the data from the source drive to the target drive.

Example

```
ufsdump 0f - <source_boot_device> | (cd /<mount_point_of_bootable_lun>;
ufsrestore rf -)
```

Note: If your configuration boots off more than one device, you must create and configure a bootable LUN that matches each boot device on the host. If you are copying non-root boot partitions, see the Oracle Corporation document [Sun StorEdge SAN Foundation Software 4.4 Configuration Guide](#).

The following example copies the information from c0t0d0s0.

```
# ufsdump 0f - /dev/rdisk/c0t0d0s0 | (cd /mnt/bootlun; ufsrestore rf -)
```

4. Use the `vi` command to edit the `/etc/vfstab` file on the SAN LUN file system.

Change the swap, root, and other file systems you will be copying from the target disk to the SAN boot disk.

Example

```
# vi /mnt/bootlun/etc/vfstab
```

The following example shows the `vfstab` entry for the SAN boot LUN.

#device	device	mount	FS	fsck
#to mount	to fsck	point	type	pass at
boot options				
#				
fd	-	/dev/fd	fd	- no -
/proc	-	/proc	proc	- no -
/dev/dsk/c6t60A98000316B61386B5D425A38797065d0s1	-	-	-	-
- swap	-	-	-	-
/dev/dsk/c6t60A98000316B61386B5D425A38797065d0s0	-	/dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0	/	ufs 1
no	-	-	-	-
/dev/dsk/c6t60A98000316B61386B5D425A38797065d0s6	-	/dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s6	/globaldevices	ufs
2	yes	-	-	-
/devices	-	/devices	devfs	- no -
sharefs	-	/etc/dfs/sharetab	sharefs	- no -
ctfs	-	/system/contract	ctfs	- no -
objfs	-	/system/object	objfs	- no -
swap	-	/tmp	tmpfs	- yes -

5. Unmount the file system on the SAN Boot LUN.

Example

```
# umount /mnt/bootlun
```

6. Repeat steps 1-3 for each file system that needs to be copied to the SAN boot LUN.

You do not need to repeat step 4 since that is only done on the root slice.

7. After the data has been copied, you will need to install the boot block or GRUB boot loader.

Related concepts

[Making the SAN boot LUN bootable](#) on page 102

ZFS File systems: Copying data from locally booted disk

The steps to copy the root zpool to the SAN boot LUN differ depending on whether you are using Solaris 10 or Solaris 11. Solaris 10 uses the Live Upgrade tools and Solaris 11 uses the beadm tools.

Note: If you have ZFS file systems in your root zpool with alternate mount points that do not use the form of `/<zpool>/<filesystem>`, you will need to take extra steps to copy them to the SAN boot zpool.

Your root zpool and SAN boot zpool cannot use the same mount point. Otherwise, you may have a failure during boot. Move any file systems on the source zpool that use the alternate mount point option to "legacy" mode before rebooting to the SAN boot LUN.

Related tasks

[Solaris 10 ZFS: Copying data from a locally booted disk](#) on page 98

[Solaris 11 ZFS: Copying data from a locally booted disk](#) on page 100

Solaris 10 ZFS: Copying data from a locally booted disk

For Solaris 10 hosts booted off of ZFS root pools, the system will be pre-configured with a Live Upgrade boot environment. To deploy a SAN booted LUN inside a root zpool, you can create and activate a new boot environment.

About this task

In the example below, the new boot environment is called Solaris10 SANBOOT and the new zpool is called `rpool_san`.

At the time of this publication, there is an issue for Solaris X64 systems, where `installgrub` fails to install the boot loader on disks that do not use 512-byte aligned LUNs. Solaris Host Utilities 6.1 defines the block size for IBM LUNs as 4,096 (4K). If you want to SAN boot an X64 host there is a work around using LOFI devices to build the zpool. The `-P` option to `lofiadm` is a hidden option that is available in Solaris 10 when kernel patch 147441-19 or later is loaded on your X64 host. These steps are documented below.

Steps

1. Create the new zpool using slice 0.

- SPARC Hosts:


```
zpool create <sanboot_pool> <disk>s0
```

```
# zpool create rpool_san c6t60A98000316B61386B5D425A38797065d0s0
```

- X86 Hosts:

```
# lofiadm -a /dev/dsk/c6t60A98000316B61386B5D425A38797065d0s0
# lofiadm -P 512 /dev/dsk/c6t60A98000316B61386B5D425A38797065d0s0
# zpool create rpool_san /dev/lofi/1
# zpool export rpool_san
# lofiadm -d /dev/lofi/1
# zpool import rpool_san
```

2. Create the new boot environment on the new pool.

Example

```
lucreate -n <new_name> -p <new_zpool>
```

```
# lucreate -n Solaris10_SANBOOT -p rpool_san
```

3. Temporarily mount the new boot environment.

Example

```
# lumount Solaris10_SANBOOT /mnt
```

4. Verify that the swap entry in the vfstab file has been set up correctly on the new boot environment.

Example

```
# vi /mnt/etc/vfstab
```

Look for the following output and make sure the zpool matches the new boot environment's zpool. If it is commented out or does not exist, you will need to fix the line. Pay special attention to the name of the zpool.

```
/dev/zvol/dsk/rpool_san/swap      -      -      swap      -
no      -
```

5. Verify that the dumpadm configuration file /etc/dumpadm.conf has been set up correctly on the boot LUN.

Example

```
# vi /mnt/etc/dumpadm.conf
```

Look for the following output and make sure the zpool matches the new boot environment's zpool.

```
DUMPADM_DEVICE=/dev/zvol/dsk/rpool_san/dump
```

6. Unmount the boot environment.

Example

```
#luumount Solaris 10_SANBOOT
```

Related tasks

[Veritas DMP Systems: Gathering SAN boot LUN information](#) on page 87

[Native MPxIO Systems: Gathering SAN boot LUN information](#) on page 90

[Gathering source disk information](#) on page 92

Solaris 11 ZFS: Copying data from a locally booted disk

or Solaris 11 hosts booted off of ZFS root pools, the system will be pre-configured with a boot environment. To deploy a SAN booted LUN inside a root zpool, you create a new boot environment and then activate it.

Before you begin

You should have:

- Created slice 0.

About this task

In the example below, the new boot environment is called Solaris11_SANBOOT and the new zpool is called rpool_san.

At the time of this publication there is an issue for Solaris X64 systems where installgrub fails to install the boot loader on disks that do not use 512-byte aligned LUNs. Solaris Host Utilities 6.1 defines the block size for IBM LUNs as 4,096 (4K). If you want to SAN boot an X64 host, there is a workaround using LOFI devices to build the zpool. These steps are documented below.

Steps

1. Create the new zpool using slice 0.

- SPARC Hosts:


```
zpool create <sanboot_pool> <disk>s0
```

```
# zpool create rpool_san c6t60A98000316B61386B5D425A38797065d0s0
```

- X86 Hosts:

```
# lofiadm -a /dev/dsk/c6t60A98000316B61386B5D425A38797065d0s0
# lofiadm -P 512 /dev/dsk/c6t60A98000316B61386B5D425A38797065d0s0
```

```
# zpool create rpool_san /dev/lofi/1
# zpool export rpool_san
# lofiadm -d /dev/lofi/1
# zpool import rpool_san
```

2. Use the `beadm` command to create the new boot environment on the new zpool.

Example

```
beadm create -p <new_pool> <new_boot_environment_name>
```

```
# beadm create -p rpool_san Solaris11_SANBOOT
```

3. Use the `zfs create` command to create an entry for swap.

The swap volume needs to be a zvol.

Example

In this example, 4g is used for swap. Choose size that is appropriate for your environment.

```
zfs create -V <size> <new_pool>/swap
```

```
# zfs create -V 4g rpool_san/swap
```

4. Use the `zfs create` command to create an entry for dump.

The dump volume needs to be a zvol.

Example

In this example, 5g is used for dump. Choose size that is appropriate for your environment.

```
zfs create -V <size> <new_pool>/dump
```

```
# zfs create -V 5g rpool_san/dump
```

5. If there are any additional ZFS file systems that exist in your source zpool, but that were not part of the boot environment definition, you will need to create them now on the new zpool.
6. Temporarily mount the new boot environment.

Example

```
# beadm mount Solaris11_SANBOOT /mnt
```

7. Use the `vi` command to verify that the swap entry was set up correctly in the `vfstab` file on the new boot environment.

Example

```
# vi /mnt/etc/vfstab
```

Look for the following output and make sure the zpool matches the new boot environment zpool. If it is commented out or does not exist, you will need to fix the line. Pay special attention to the name of the zpool.

```
/dev/zvol/dsk/rpool_san/swap      -      -      swap      -
no      -
```

8. Use the `vi` command to verify that the `dupadm` configuration file `/etc/dumpadm.conf` has been set up correctly on the boot LUN.

Example

```
# vi /mnt/etc/dumpadm.conf
```

Look for the following output and make sure the zpool matches the new boot environment's zpool.

```
DUMPADM_DEVICE=/dev/zvol/dsk/rpool_san/dump
```

9. Unmount the boot environment.

Example

```
# beadm unmount Solaris11_SANBOOT
```

Related tasks

[Veritas DMP Systems: Gathering SAN boot LUN information](#) on page 87

[Native MPxIO Systems: Gathering SAN boot LUN information](#) on page 90

[Gathering source disk information](#) on page 92

Making the SAN boot LUN bootable

The process of making the SAN boot LUN bootable differs depending on whether you are using a SPARC system or an X64 system. For SPARC systems, a boot block must be installed. For X64 systems, the GRUB bootloader must be installed to the disk.

SPARC: Installing the boot block

For SPARC systems, the `installboot` command is used to install the boot block to a disk.

Before you begin

You will need to use the `/dev/rdisk` path to the SAN boot LUN.

Solaris 10 UFS Hosts

Step

1. Run the `installboot` command.

Example

```
/usr/sbin/installboot -F ufs /usr/platform/`uname -i`/lib/fs/ufs/  
bootblk /dev/rdisk/<bootlun>s0
```

```
# /usr/sbin/installboot -F ufs /usr/platform/`uname -i`/lib/fs/ufs/  
bootblk  
/dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0
```

Solaris 10 ZFS Hosts

For Solaris 10 ZFS SPARC hosts, you do not need to install the boot block manually. The boot block will be installed during the boot environment activation process.

About this task

Step

1. Activate the boot environment.

You should have created the boot environment earlier.

Example

In this example, the new boot environment is named `Solaris10_SANBOOT`.

```
luactivate <sanboot_boot_environment>
```

```
# luactivate Solaris10_SANBOOT
```

Solaris 11 ZFS Hosts

For Solaris 11 ZFS SPARC hosts, you need to install the boot block after activating the boot environment.

Steps

1. Use the `beadm` command to activate the boot environment.

Example

In this example, the new boot environment is named Solaris11_SANBOOT.

```
beadm activate <sanboot_boot_environment>
```

```
# beadm activate Solaris11_SANBOOT
```

Note: The `beadm activate` command does not actually install the boot loader on the SAN boot LUN when the boot environment resides in a different zpool.

2. Install the boot block on the SAN boot LUN.

Example

```
/usr/sbin/installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/  
bootblk /dev/rdisk/<bootlun>s0
```

```
# /usr/sbin/installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/  
bootblk /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0
```

X64: Installing GRUB information

For X64 systems, the `installgrub` command is used to install the GRUB boot loader to a disk.

Before you begin

You will need the `/dev/rdisk` path to the SAN boot LUN. You will also need to update the `bootenv.rc` file on Solaris 10 hosts to indicate the proper bootpath.

Solaris 10 UFS Hosts

For Solaris 10 X64 UFS hosts, you need to update the `bootenv.rc` file to set the bootpath, update the boot archive, and install the GRUB boot loader.

Steps

1. Mount the file system that you will use as your target boot device.

Example

```
# mount /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0 /mnt/bootlun
```

2. Use the `vi` command to edit the `/mnt/bootlun/boot/solaris/bootenv.rc` file.

Example

```
# vi /mnt/bootlun/boot/solaris/bootenv.rc
```

3. Find the output line that starts with "setprop bootpath" and adjust the device path to that of the HBA, WWPN, and hex address of the SAN boot LUN.

Example

Change:

```
setprop bootpath '/pci@0,0/pci8086,25f8@4/pci1000,3150@0/sd@5,0:a'
```

To:

```
setprop bootpath '/pci@0,0/pci8086,25f9@6/pci10df,fe00@0/fp@0,0/disk@w500a0983974988c3,z:a'
```

4. Update the boot archive on the SAN boot LUN

Example

```
# bootadm update-archive -R /mnt/bootlun
```

5. Run the installgrub command.

Example

```
cd /boot/grub; /sbin/installgrub stage1 stage2 /dev/rdisk/<bootlun>s0
```

```
# cd /boot/grub; /sbin/installgrub stage1 stage2 /dev/rdisk/c6t60A98000316B61386B5D425A38797065d0s0
```

6. Unmount the SAN boot LUN.

Example

```
# umount /mnt/bootlun
```

Solaris 10 ZFS Hosts

For Solaris 10 ZFS X64 hosts, you do not need to install the GRUB information manually. The GRUB boot loader will be installed during the boot environment activation process.

Step

1. Activate the boot environment which you created earlier.

Example

In this example, the new boot environment is named: Solaris10_SANBOOT.

```
luactivate <sanboot_boot_environment>
```

```
# luactivate Solaris10_SANBOOT
```

Solaris 11 ZFS Hosts

For Solaris 11 ZFS X64 hosts, we need to install the GRUB boot loader after activating the boot environment.

Steps

1. Activate the boot environment using the `beadm activate` command.

Example

In this example, the boot environment is `Solaris11_SANBOOT`.

```
beadm activate <sanboot_boot_environment>
```

```
# beadm activate Solaris11_SANBOOT
```

Note: The `beadm activate` command doesn't actually install the boot loader on the SAN boot LUN when the boot environment resides in a different zpool.

2. (Optional): If you are using a headless server and only using a console connection, you will need to modify the GRUB `menu.st` file on the SAN boot zpool and disable the splashimage entry.

Example

```
vi /<sanboot_zpool>/boot/grub/menu.lst
```

```
# vi /rpool_san/boot/grub/menu.lst
```

- a) Look for this line: `#splashimage /boot/grub/splash.xpm.gz`
- b) Comment it out.

Example

```
#splashimage /boot/grub/splash.xpm.gz
```

3. Install the boot loader on the SAN boot LUN.

Example

In this example, the boot environment is `Solaris11_SANBOOT` and the zpool is `rpool_san`.

```
# bootadm install-bootloader -P <sanboot_zpool>
```

```
# bootadm install-bootloader -P rpool_san
```

Configuring the host to boot from the SAN boot LUN

The procedures for configuring the host to boot from the SAN boot LUN differ depending on whether your system is a SPARC host or an X64 host.

Configuring the host to boot from the SAN boot LUN on SPARC-based systems

In order to configure the SPARC-based host to boot from the SAN-based LUN, you need to shut down the host to the OpenBoot Prompt, create aliases for the device path, and boot up the system.

Before you begin

Be sure that the WWPN and target used match the device used in the `vfstab` for UFS systems. If you do not use the correct path and device, the system will not fully boot.

Steps

1. Halt the host by running `init 0` or `shutdown -i 0 -g 0 -y`

Example

```
# sync; sync; init 0 or
```

```
or
```

```
# shutdown -i 0 -g 0 -y
```

2. Run `show-disks` to list the device paths to your disks.

Example

```
ok> show-disks
a) /pci@7c0/pci@0/pci@9/emlx@0,1/fp@0,0/disk
b) /pci@7c0/pci@0/pci@9/emlx@0/fp@0,0/disk
c) /pci@7c0/pci@0/pci@1/pci@0/ide@8/cdrom
d) /pci@7c0/pci@0/pci@1/pci@0/ide@8/disk
e) /pci@780/pci@0/pci@9/scsi@0/disk
g) NO SELECTION
Enter Selection, q to quit:
```

3. Find the path that you want to use with the boot device.

This path should match the path you displayed and documented earlier when using the `luxadm display` command. Make sure you match the WWPN of the target with the correct HBA device path.

4. Create the device alias for your SAN boot device.

Example

This example uses: `"/pci@7c0/pci@0/pci@9/emlx@0,1/fp@0,0/disk"` and WWPN, LUN `"500a0984974988c3,d3"`

```
ok nvalias sanboot /pci@7c0/pci@0/pci@9/emlx@0,1/fp@0,0/
disk@w500a0984974988c3,d3:a
ok nvstore
```

5. Set the boot-device to the new device alias.

Example

```
ok setenv boot-device sanboot
```

6. Boot the system.

Example

```
ok boot
```

Related tasks

[Veritas DMP Systems: Gathering SAN boot LUN information](#) on page 87

[Native MPxIO Systems: Gathering SAN boot LUN information](#) on page 90

[Gathering source disk information](#) on page 92

Configuring the host to boot from the SAN boot LUN on X64-based systems

In order to configure the X64-based host to boot from the SAN-based LUN, you need to configure the HBA boot bios to boot the LUN. You may also need to configure the system BIOS to boot from the HBA.

Before you begin

You will need the information gathered previously when configuring the bios. Also, be sure the WWPN and target used match the device used in the `vfstab` for UFS systems. If you do not use the correct path and device, the system will not fully boot.

About this task

The configuration of the HBA boot bios be done using the QLogic or Emulex HBA utilities while the operating system is running or while the system is booting.

Steps

1. Emulex HBA: Use the `hbacmd` Emulex utility to configure the LUN and target WWPN used by the HBA boot bios.

You can skip this step if you would prefer to set the information during the POST process of the system.

Note: You can only set one option at a time with `hbacmd` so you will need to run the command once to set the LUN and once to set the target wwpn.

Example

```
hbacmd setbootparams <initiator_wwpn> x86 LUN <LUN_DECIMAL> BootDev 0
```

```
# hbacmd setbootparams 10:00:00:00:c9:6e:98:55 x86 LUN 213 BootDev 0
```

```
hbacmd setbootparams <initiator_wwpn> x86 TargetWwpn <Target_WWPN>  
BootDev 0
```

```
# hbacmd setbootparams 10:00:00:00:c9:6e:98:55 x86 TargetWwpn 50:0a:  
09:84:97:49:88:c3 BootDev 0
```

2. Qlogic HBA: Use the `scli` utility to configure the LUN and target used by the HBA boot BIOS

See the vendor documentation for specifics on how to use these commands because the exact steps may vary from release to release. You can skip this step if you would prefer to set the information during the POST process of the system.

3. Use one of the following commands to reboot the host.

- Solaris 10:

```
# sync; sync; init 6
```

 or

```
# shutdown -i 6 -g 0 -y
```
- Solaris 11:

```
# reboot -p
```

Note: Solaris 11 now defaults to a fast reboot that bypasses system POST.

4. (Optional): Configure the HBA during system POST to boot from the SAN boot LUN.

The system will restart the POST process. You can skip this step if you configured the boot BIOS while the system was still booted.

5. Enter setup mode for the system BIOS and configure the boot device order.

Depending on the system type and BIOS type, you might see an item called Hard Disks. Hard Disks will have a submenu that lists all of the possible hard disks available to boot from. Make sure the SAN boot LUN is first in the list.

6. Boot the system

Related tasks

Veritas DMP Systems: Gathering SAN boot LUN information on page 87

Native MPxIO Systems: Gathering SAN boot LUN information on page 90

Gathering source disk information on page 92

Veritas DMP: Enabling root encapsulation

For Veritas DMP systems, you should encapsulate the root disk so that Veritas DMP multipathing is enabled.

About this task

To encapsulate the root disk, run `vxdiskadm` and choose menu option 2. Your host will reboot on the newly encapsulated volume. The `/etc/vfstab` will be updated during the encapsulation process to use the new encapsulated volume

Steps

1. Run `vxdiskadm`.
2. Choose option 2.
3. Follow the directions presented by the `vxdiskadm` command.
4. Reboot your host.

Supported Solaris and Data ONTAP features

The Host Utilities work with both Solaris and Data ONTAP features.

Features supported by the Host Utilities

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

Some of the supported features include:

- Multiple paths to the storage system when a multipathing solution is installed
- Veritas VxVM, Solaris Volume Manager (SVM), and ZFS file systems
- **(MPxIO)** ALUA
- Oracle VM Server for SPARC (Logical Domains)
- SAN booting

For information on which features are supported with which configurations, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

Related information

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

HBAs and the Solaris Host Utilities

The Host Utilities support a number of HBAs.

Ensure the supported HBAs are installed before you install the Host Utilities. Normally, the HBAs should have the correct firmware and FCode set. To determine the firmware and FCode setting on your system, run the appropriate administration tool for your HBA.

Note: For details on the specific HBAs that are supported and the required firmware and FCode values, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

Related information

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

Multipathing and the Solaris Host Utilities

The Solaris 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 MPxIO environment of the Host Utilities uses Oracle's native multipathing solution (MPxIO).

Note: The Host Utilities also support IP Multipathing (IPMP). You do not need to perform any specific N series configuration to enable IPMP.

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

iSCSI and multipathing

You can use iSCSI with either Veritas DMP or MPxIO.

You should have at least two Ethernet interfaces on the storage system enabled for iSCSI traffic. Having two interfaces enables you to take advantage of multipathing. Each iSCSI interface must be in a different iSCSI target portal group.

In Veritas DMP environments, you must also disable MPxIO before you can use iSCSI. You must use DMP for multipathing when you are using Veritas.

For more information about using multipathing with iSCSI, see *Using iSCSI Multipathing in the Solaris 10 Operating System*.

Volume managers and the Solaris Host Utilities

The Solaris Host Utilities support different volume management solutions based on your environment.

The Veritas DMP environment uses Veritas Volume Manager (VxVM).

The MPxIO stack works with Solaris Volume Manager (SVM), ZFS, and VxVM to enable you to have different volume management solutions.

Note: To determine which versions of VxVM are supported with MPxIO, see the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Related information

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

(FC) ALUA support with certain versions of Data ONTAP

The MPxIO environment of the Solaris Host Utilities requires that you have ALUA enabled for high availability storage controllers (clustered storage systems) using FC and a version of Data ONTAP that supports ALUA. Veritas Storage Foundation also supports ALUA starting with version 5.1 P1.

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

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

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

Note: iSCSI is not supported with ALUA if you are running Data ONTAP operating in 7-Mode or Data ONTAP before 8.1.1. operating in Cluster-Mode.

Check your version of Data ONTAP to see if it supports ALUA and check the N series support website (accessed and navigated as described in [Websites](#) on page 9) to see if the Host Utilities support that version of Data ONTAP. ALUA support was first available with Data ONTAP 7.2 and single-image cfmode.

You can also check the N series support website (accessed and navigated as described in [Websites](#) on page 9) to determine if your version of the Host Utilities supports Veritas Storage Foundation 5.1 P1 or later.

(FC) Solaris Host Utilities configurations that support ALUA

The Solaris Host Utilities support ALUA in both MPxIO environments and certain Veritas Storage Foundation environments as long as the environments are running the FC protocol. ALUA is only supported in environments running the iSCSI protocol with Clustered ONTAP.

If you are using MPxIO with FC and high availability storage controllers with any of the following configurations, you must have ALUA enabled:

Host Utilities version	Host requirements	Data ONTAP version
Host Utilities 6.1	Solaris 10u8 and later	7.3.5.1 and later
Host Utilities 6.0	Solaris 10 update 5 and later	7.3 and later
Host Utilities 4.1 through 5.1	Solaris 10 update 3 and later	7.2.1.1 and later
Host Utilities 4.0	Solaris 10 update 2 only with QLogic drivers and SPARC processors	7.2.1 and later

If you are running the Host Utilities with Veritas Storage Foundation 5.1 P1 and the FC protocol, you can use ALUA.

Note: Earlier versions of Veritas Storage Foundation do not support ALUA.

Oracle VM Server for SPARC (Logical Domains) and the Host Utilities

Certain configurations of the Host Utilities MPxIO stack support Oracle VM Server for SPARC (Logical Domains).

The supported configurations include guests that are I/O Domains or guests that have iSCSI configured. You must install the Host Utilities if a guest is using N series storage.

If you are using Logical Domains (Oracle VM Server for SPARC), you must configure your system with the Host Utilities settings. You can use Host Utilities `host_config` command do this or you can configure the settings manually.

A Solaris host running Logical Domains accesses and uses N series storage exactly the same way any other Solaris host does.

For information on which configurations support Logical Domains, see the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Related information

The N series support website: www.ibm.com/storage/support/nseries/

SAN booting and the Host Utilities

The Host Utilities support SAN booting with both the Veritas DMP and MPxIO environments. SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for a Solaris 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.

Support for non-English versions of Solaris operating systems

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

High-level look at Host Utilities Veritas DMP stack

The Host Utilities Veritas DMP stack works with Solaris hosts running Veritas Storage Foundation.

The following is a high-level summary of the supported Veritas DMP stack.

Note: Check the N series support website (accessed and navigated as described in [Websites](#) on page 9) for details and current information about the supported stack.

- Operating system:
 - Solaris 10 update 8 and later. See the N series support website (accessed and navigated as described in [Websites](#) on page 9) for more information.
- Processor:
 - SPARC processor systems
 - x86/64 processor systems
- FC HBA
 - Emulex LPFC HBAs and their Oracle-branded equivalents
 - Certain Oracle OEM QLogic® HBAs and their Oracle-branded equivalents
 - Certain Oracle OEM Emulex® HBAs and their Oracle-branded equivalents
- iSCSI software initiators
- Drivers
 - Oracle-branded Emulex drivers (emlxs)
 - Oracle-branded QLogic drivers (qlc)
- Multipathing
 - Veritas DMP

The Host Utilities Veritas DMP stack also supports the following:

- Volume manager
 - VxVM
- Clustering
 - Veritas Cluster Server (VCS)

Related information

The N series support website: www.ibm.com/storage/support/nseries/

High-level look at Host Utilities MPxIO stack

The Host Utilities MPxIO stack works with Solaris hosts running Oracle StorEdge SAN Foundation Software and components that make up the native stack.

The following is a high-level summary of the supported MPxIO stack at the time this document was produced.

Note: Check the N series support website (accessed and navigated as described in [Websites](#) on page 9) for details and current information about the supported stack.

- Operating system:
 - Solaris 10 update 8 and later
- Processor:
 - SPARC processor systems
 - x86/64 processor systems
- HBA
 - Certain QLogic HBAs and their Oracle-branded equivalents
 - Certain Emulex HBAs and their Oracle-branded equivalents
- Drivers
 - Bundled Oracle StorEdge SAN Foundation Software Emulex drivers (emlxs)
 - Bundled Oracle StorEdge San Foundation Software QLogic drivers (qlc)
- Multipathing
 - Oracle StorageTek Traffic Manager (MPxIO)

The Host Utilities MPxIO stack also supports the following:

- Volume manager
 - SVM
 - VxVM
 - ZFS

Note: To determine which versions of VxVM are supported with MPxIO, see the Interoperability Matrix.
- Clustering
 - Sun Clusters. This kit has been certified using the Sun Cluster Automated Test Environment (SCATE)
 - Veritas Cluster Server (VCS)

Related information

The N series support website: www.ibm.com/storage/support/nseries/

Protocols and configurations supported by the Solaris Host Utilities

The Solaris Host Utilities provide support for FC and iSCSI connections to the storage system using direct-attached, fabric-attached, and network configurations.

Notes about the supported protocols

The FC and iSCSI protocols enable the host to access data on storage systems.

The storage systems are targets that have storage target devices called LUNs. The FC and iSCSI protocols enable the host to access the LUNs to store and retrieve data.

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

The FC protocol

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

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

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

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

The iSCSI protocol

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

On the host, the iSCSI protocol is implemented over either the host's standard Ethernet interfaces or on an HBA.

On the storage system, the iSCSI protocol can be implemented over the storage system's standard Ethernet interface using one of the following:

- A software driver that is integrated into Data ONTAP

- (Data ONTAP 7.1 and later) An iSCSI target HBA or an iSCSI TCP/IP offload engine (TOE) adapter. You do not have to have a hardware HBA.

The connection between the host and storage system uses a standard TCP/IP network. The storage system listens for iSCSI connections on TCP port 3260.

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

Supported configurations

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

The Host Utilities support the following basic configurations:

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

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

- FC direct-attached. A single host with a direct FC connection from the HBA to stand-alone or active/active storage system configurations.
- iSCSI direct-attached. One or more hosts with a direct iSCSI connection to stand-alone or active/active storage systems. The number of hosts that can be directly connected to a storage system or a pair of storage systems depends on the number of available Ethernet ports.
- iSCSI network-attached. In an iSCSI environment, all methods of connecting Ethernet switches to a network that have been approved by the switch vendor are supported. Ethernet switch counts are not a limitation in Ethernet iSCSI topologies. Refer to the Ethernet switch vendor documentation for specific recommendations and best practices.

The *SAN Configuration Guide* (called *Fibre Channel and iSCSI Configuration Guide* in Data ONTAP 8.1 and earlier) provides detailed information, including diagrams, about the supported topologies. There is also configuration information in the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP. Refer to those documents for complete information about configurations and topologies.

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.

About the troubleshooting sections that follow

The troubleshooting sections that follow help you verify your system setup.

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 about Solaris commands, see the man pages and operating system documentation.
- For more information about the Data ONTAP commands, see the Data ONTAP documentation, in particular, the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP.
- 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 can 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 Interoperability Matrix. Support for new components is added on an ongoing basis. This online document contains a complete list of supported HBAs, platforms, applications, and drivers.

Related information

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

Check the version of your host operating system

Make sure you have the correct version of the operating system.

You can use the `cat /etc/release` command to display information about your operating system.

The following example checks the operating system version on a SPARC system.

```
# cat /etc/release
                Solaris 10 5/08 s10s_u5wos_10 SPARC
Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.
                Use is subject to license terms.
                Assembled 24 March 2008
```

The following example checks the operating system version on an x86 system.

```
# cat /etc/release
                Solaris 10 5/08 s10x_u5wos_10 X86
Copyright 2008 Sun Microsystems, Inc. All Rights Reserved.
                Use is subject to license terms.
                Assembled 24 March 2008
```

Confirm the HBA is supported

You can use the `sanlun` command to display information on the HBA and the Interoperability Matrix to determine if the HBA is supported. Supported HBAs should be installed before you install the Host Utilities.

The `sanlun` command is part of the Host Utilities SAN Toolkit.

If you are using MPxIO, you can also use the `fcinfo hba-port` command to get information about the HBA.

1. The following example uses the `sanlun` command to check a QLogic HBA in an environment using a Solaris native qlc driver.

```
sanlun fcp show adapter -v

adapter name:      qlc1
WWPN:              210000e08b88b838
WWNN:              200000e08b88b838
driver name:       20060630-2.16
model:             QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:     Not Available
hardware version:  Not Available
driver version:    20060630-2.16
firmware version:  4.0.23
Number of ports:   1 of 2
port type:         Fabric
port state:        Operational
supported speed:   1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed:  4 GBit/sec
OS device name:    /dev/cfg/c2
```

```

adapter name:      qlc2
WWPN:             210100e08ba8b838
WWNN:             200100e08ba8b838
driver name:      20060630-2.16
model:            QLA2462
model description: Qlogic PCI-X 2.0 to 4Gb FC, Dual Channel
serial number:    Not Available
hardware version: Not Available
driver version:   20060630-2.16
firmware version: 4.0.23
Number of ports: 2 of 2
port type:        Fabric
port state:       Operational
supported speed:  1 GBit/sec, 2 GBit/sec, 4 GBit/sec
negotiated speed: 4 GBit/sec
OS device name:   /dev/cfg/c3

```

Related information

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

(MPxIO, native drivers) Ensure that MPxIO is configured correctly for ALUA on FC systems

Configurations using native drivers with MPxIO and FC in a clustered environment require that ALUA be enabled.

In some cases, ALUA might have been disabled on your system and you will need to re-enable it. For example, if your system was used for iSCSI or was part of a single storage controller configuration, the `symmetric-option` might have been set. This option disables ALUA on the host.

To enable ALUA, you must remove the `symmetric-option` by doing one of the following:

- Running the `host_config` command. This command automatically comments out the `symmetric-option` section.
- Editing the appropriate section in the `/kernel/drv/scsi_vhci.conf` for Solaris 10 or `/etc/driver/drv/scsi_vhci.conf` file for Solaris 11 to manually comment it out. The example below displays the section you must comment out.

Once you comment out the option, you must reboot your system for the change to take effect.

The following example is the section of the `/kernel/drv/scsi_vhci.conf` file for Solaris 10 or `/etc/driver/drv/scsi_vhci.conf` file for Solaris 11 that you must comment out if you want to enable MPxIO to work with ALUA. This section has been commented out.

```
#device-type-iscsi-options-list =
#"NETAPP LUN", "symmetric-option";
#symmetric-option = 0x1000000;
```

Ensure that MPxIO is enabled on SPARC systems

When you use a MPxIO stack on a SPARC system, you must manually enable MPxIO. If you encounter a problem, make sure that MPxIO is enabled.

Note: On x86/64 systems, MPxIO is enabled by default.

To enable MPxIO on a SPARC system, use the `stmsboot` command. This command modifies the `fp.conf` file to set the `mpxio_disable=` option to `no` and updates `/etc/vfstab`.

After you use this command, you must reboot your system.

The options you use with this command vary depending on your version of Solaris. For systems running Solaris 10 update 5, execute: `stmsboot -D fp -e`

For example, if MPxIO is not enabled on a system running Solaris 10 update 5, you would enter the following commands. The first command enables MPxIO by changing the `fp.conf` file to read `mpxio_disable=no`. It also updates `/etc/vfstab`. You must reboot the system for the change to take effect. Input the following commands to reboot the system.

For FC:

```
# stmsboot -D fp -e
# touch /reconfigure
# init 6
```

For iSCSI

```
# stmsboot -D fp -e
# touch /reconfigure
# init 6
```

(MPxIO) Ensure that MPxIO is enabled on iSCSI systems

While MPxIO should be enabled by default on iSCSI systems, you can confirm this by viewing the iSCSI configuration file `/kernel/drv/iscsi.conf` file.

When MPxIO is enabled, this file has the `mpxio-disable` set to `"no"`.

```
mpxio-disable="no"
```

If this line is set to `"yes"`, you must change it by doing one of the following:

- Running the `host_config` command. This command sets the symmetric option.
- Editing the appropriate section in the `/kernel/drv/iscsi.conf` file for Solaris 10 and `/etc/driver/drv/iscsi.conf` file for Solaris 11 to manually set the command to "no". The example below displays the section you must comment out.

You **must** reboot your system for the change to take effect.

Here is an example of a `/kernel/drv/iscsi.conf` file for Solaris 10 that has MPxIO enabled. The line that enables MPxIO, `mpxio-disable="no"` is in bold to make it easier to locate.

```
# Copyright 2006 Sun Microsystems, Inc. All rights reserved.
# Use is subject to license terms.
#
#ident "@(#)iscsi.conf 1.2 06/06/12 SMI"
name="iscsi" parent="/" instance=0;
ddi-forceattach=1;
Chapter 3: Configuring the initiator 23
#
# I/O multipathing feature (MPxIO) can be enabled or disabled using
# mpxio-disable property. Setting mpxio-disable="no" will activate
# I/O multipathing; setting mpxio-disable="yes" disables the
# feature.
#
# Global mpxio-disable property:
#
# To globally enable MPxIO on all iscsi ports set:
# mpxio-disable="no";
#
# To globally disable MPxIO on all iscsi ports set:
# mpxio-disable="yes";
#
mpxio-disable="no";
tcp-nodelay=1;
...
```

(MPxIO) Verify that MPxIO multipathing is working

You can confirm that multipathing is working in an MPxIO environment by using either a Host Utilities tool such as the `sanlun lun show` command or a Solaris tool such as the `mpathadm` command.

The `sanlun lun show` command displays the disk name. If MPxIO is working, you should see a long name similar to the following:

```
/dev/rdisk/c5t60A980004334686568343771474A4D42d0s2
```

The long, consolidated Solaris device name is generated using the LUN's serial number in the IEEE registered extended format, type 6. The Solaris host receives this information in the SCSI Inquiry response.

Another way to confirm that MPxIO is working is to check for multiple paths. To view the path information, you need to use the `mpathadm` command. The `sanlun` cannot display the underlying multiple paths because MPxIO makes these paths transparent to the user when it displays the consolidated device name shown above.

In this example, the `mpathadm list lu` command displays a list of all the LUNs.

```
# mpathadm list lu
  /dev/rdisk/c3t60A980004334612F466F4C6B72483362d0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A980004334612F466F4C6B72483230d0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A980004334612F466F4C6B7248304Dd0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A980004334612F466F4C6B7247796Cd0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A980004334612F466F4C6B72477838d0s2
    Total Path Count: 8
    Operational Path Count: 8
  /dev/rdisk/c3t60A980004334612F466F4C6B72477657d0s2
    Total Path Count: 8
    Operational Path Count: 8
```

(Veritas DMP) Check that the ASL and APM have been installed

You must have the ASL and APM installed in order for the Veritas DMP to identify whether the path is primary or secondary

Without the ASL and APM, the DSM treats all paths as equal, even if they are secondary paths. As a result, you might see I/O errors on the host. On the storage system, you might see the Data ONTAP error:

```
FCP Partner Path Misconfigured
```

If you encounter the I/O errors on the host or the Data ONTAP error, make sure you have the ASL and APM installed.

(Veritas) Check VxVM

The Host Utilities support the VxVM for the Veritas DMP stack and certain configurations of the MPxIO stack. You can use the `vxdisk list` command to quickly check the VxVM disks and the `vxprint` command to view volume information.

Note: (MPxIO) To determine which versions of VxVM are supported with MPxIO, see the N series interoperability matrix website (accessed and navigated as described in [Websites](#) on page 9).

See your Veritas documentation for more information on working with the VxVM.

Related information

The N series support website: www.ibm.com/storage/support/nseries/

(MPxIO) Check the Solaris Volume Manager

If you are using the MPxIO version of the Host Utilities with the Solaris Volume Manager (SVM), it is a good practice to check the condition of the volumes.

The `metastat -a` command lets you quickly check the condition of SVM volumes.

Note: In a Solaris Cluster environment, metaset and their volumes are only displayed on the node that is controlling the storage.

See your Solaris documentation for more information on working with the SVM.

The following sample command line checks the condition of SVM volumes:

```
# metastat -a
```

(MPxIO) Check settings in `ssd.conf` or `sd.conf`

Verify that you have the correct settings in the configuration file for your system.

The file you need to modify depends on the processor your system uses:

- SPARC systems with MPxIO enabled use the `ssd.conf` file. You can use the `host_config` command to update the `/kernel/drv/ssd.conf` file.
- x86/64 systems with MPxIO enabled use the `sd.conf` file. You can use the `host_config` command to update the `sd.conf` file.

Example of `ssd.conf` file (MPxIO on a SPARC system):

You can confirm that the `ssd.conf` file was correctly set up by checking to ensure that it contains the following:

```
ssd-config-list="NETAPP LUN", "physical-block-size:4096,
retries-busy:30, retries-reset:30, retries-notready:300,
retries-timeout:10, throttle-max:64, throttle-min:8";
```

Example of sd.conf file (MPxIO on an x86/64 system):

You can confirm that the `sd.conf` file was correctly set up by checking to ensure that it contains the following:

```
sd-config-list="NETAPP LUN", "physical-block-size:4096,
retries-busy:30, retries-reset:30, retries-notready:300,
retries-timeout:10, throttle-max:64, throttle-min:8";
```

Check the storage system setup

Make sure your storage system is set up correctly.

(MPxIO/FC) Check the ALUA settings on the storage system

In MPxIO environments using FC, you must have ALUA set on the storage system to work with igroups.

You can verify that you have ALUA set for the igroup by executing the `igroup show -v`

Note: For more information on ALUA, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP. In particular, see the section “Enabling ALUA.”

The following command line displays information about the `cfmode` on the storage system and shows that ALUA is enabled. (To make the information on ALUA easier to locate, it is shown in bold.)

```
filerA# igroup show -v
Tester (FCP):
  OS Type: solaris
  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
```

Verifying that the switch is installed and configured

If you have a fabric-attached configuration, check that the switch is set up and configured as outlined in the instructions that shipped with your hardware.

You should have completed the following tasks:

- Installed the switch in a system cabinet or rack.
- Confirmed that the Host Utilities support this switch.
- Turned on power to the switch.
- Configured the network parameters for the switch, including its serial parameters and IP address.

Related information

The N series support website: www.ibm.com/storage/support/nseries/

Determining whether to use switch zoning

If you have a fabric-attached configuration, determine whether switch zoning is appropriate for your system setup.

Zoning requires more configuration on the switch, but it provides the following advantages:

- It simplifies configuring the host.
- It makes information more manageable. The output from the host tool `iostat` is easier to read because fewer paths are displayed.

To have a high-availability configuration, make sure that each LUN has at least one primary path and one secondary path through each switch. For example, if you have two switches, you would have a minimum of four paths per LUN.

It is recommended that your configuration have no more than eight paths per LUN. For more information about zoning, see the N series support website (accessed and navigated as described in [Websites](#) on page 9).

Power up equipment in the correct order

The different pieces of hardware communicate with each other in a prescribed order, which means that problems occur if you turn on power to the equipment in the wrong order.

Use the following order when powering on the equipment:

- Configured Fibre Channel switches
It can take several minutes for the switches to boot.
- Disk shelves
- Storage systems

- Host

Verify that the host and storage system can communicate

Once your setup is complete, make sure the host and the storage system can communicate.

You can verify that the host can communicate with the storage system by issuing a command from:

- The storage system's console
- A remote login that you established from the host

Possible iSCSI issues

The following sections describe possible issues that might occur when you are working with iSCSI.

(iSCSI) Verify the type of discovery being used

The iSCSI version of the Host Utilities supports iSNS, dynamic (SendTarget) and static discovery.

You can use the `iscsiadm` command to determine which type of discovery you have enabled.

This example uses the `iscsiadm` command to determine that dynamic discovery is being used.

```
$ iscsiadm list discovery
Discovery:
  Static: disabled
  Send Targets: enabled
  iSNS: disabled
```

(iSCSI) Bidirectional CHAP does not work

When you configure bidirectional CHAP, make sure you supply different passwords for the `inpassword` value and the `outpassword` value.

If you use the same value for both of these passwords, CHAP appears to be set up correctly, but it does not work.

(iSCSI) LUNs are not visible on the host

Storage system LUNs are not listed by the `iscsiadm list target -S` command or by the `sanlun lun show all` command.

If you encounter this problem, verify the following configuration settings:

- Network connectivity—Verify that there is TCP/IP connectivity between the host and the storage system by performing the following tasks:
 - From the storage system command line, ping the host.

- From the host command line, ping the storage system.
- Cabling—Verify that the cables between the host and the storage system are properly connected.
- System requirements—Verify that you have the correct Solaris operating system (OS) version, correct version of Data ONTAP, and other system requirements. See the appropriate interoperability matrix for your N series product, available on the IBM support website.
- Jumbo frames—If you are using jumbo frames in your configuration, ensure that jumbo frames are enabled on all devices in the network path: the host Ethernet NIC, the storage system, and any switches.
- iSCSI service status—Verify that the iSCSI service is licensed and started on the storage system. For more information about licensing iSCSI on the storage system, see the *SAN Administration Guide* (called *Block Access Management Guide for iSCSI and FC* in Data ONTAP 8.1 and earlier) for your version of Data ONTAP.
- Initiator login—Verify that the initiator is logged in to the storage system by entering the `iscsi initiator show` command on the storage system console.
If the initiator is configured and logged in to the storage system, the storage system console displays the initiator node name and the target portal group to which it is connected.
If the command output shows that no initiators are logged in, verify that the initiator is configured according to the procedure described in the section on “Configuring the initiator.”
- iSCSI node names—Verify that you are using the correct initiator node names in the `igroup` configuration.
On the storage system, use the `igroup show -v` command to display the node name of the initiator. This node name must match the initiator node name listed by the `iscsiadm list initiator-node` command on the host.
- LUN mappings—Verify that the LUNs are mapped to an `igroup` that also contains the host. On the storage system, use one of the following commands:

Data ONTAP	Command	Description
Data ONTAP operating in 7-Mode	<code>lun show -m</code>	Displays all LUNs and the <code>igroups</code> they are mapped to
Data ONTAP operating in 7-Mode	<code>lun show -g</code>	Displays the LUNs mapped to the specified <code>igroup</code>
Data ONTAP operating in Cluster-Mode	<code>lun show -m -vserver <vserver></code>	Displays the LUNs and <code>igroups</code> they are mapped to for a given Vserver

- If you are using CHAP, verify that the CHAP settings on the storage system and host match. The incoming user names and password on the host are the outgoing values on the storage system. The outgoing user names and password on the storage system are the incoming values on the host. For bidirectional CHAP, the storage system CHAP username must be set to the storage system’s iSCSI target node name.

Related information

The N series support website: www.ibm.com/storage/support/nseries

Possible MPxIO issues

The following sections describe issues that can occur when you are using the Host Utilities in an MPxIO environment.

(MPxIO) sanlun does not show all adapters

In some cases, the `sanlun lun show all` command does not display all the adapters. You can display them using either the `luxadm display` command or the `mpathadm` command.

When you use MPxIO, there are multiple paths. MPxIO controls the path over which the `sanlun` SCSI commands are sent and it uses the first one it finds. This means that the adapter name can vary each time you issue the `sanlun lun show` command.

If you want to display information on all the adapters, use either the `luxadm display` command or the `mpathadm` command. For the `luxadm display` command, you would enter

```
luxadm display -v device_name
```

Where *device_name* is the name of the device you are checking.

(MPxIO) Solaris log message says data not standards compliant

When running the Host Utilities, you might see a message in the Solaris log saying that data is not standards compliant. This message is the result of a Solaris bug.

```
WARNING: Page83 data not standards compliant
```

This erroneous Solaris log message has been reported to Oracle. The Solaris initiator implements an older version of the SCSI Spec.

The SCSI target is standards compliant, so ignore this message.

Installing the nSANity data collection program

Obtain and install the nSANity Diagnostic and Configuration Data Collector program when instructed to do so by your technical support representative.

About this task

Contact your technical support representative to obtain the nSANity Diagnostic and Configuration Data Collector program. 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.

Steps

1. Obtain the Windows zip or Linux tgz version of the nSANity program from your technical support representative and copy it to the workstation or server that you want to run it on.
2. Change to the directory to which you downloaded the zip or tgz file.
3. 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.

LUN types, OS label, and OS version combinations for achieving aligned LUNs

To align LUNs, you use certain combinations of LUN types, OS labels, and OS versions.

OS Version	Kernel	Volume Manager	OS Label	LUN Type	(s)sd.conf changes	Host Utility
11	SRU10.5	ZFS	EFI	solaris	physical-block-size: 4096	6.1
11	SRU10.5	SVM Slice	SMI	solaris		6.1
11	SRU10.5	SVM Slice	EFI	solaris_efi		6.1
10u8, 10u9, 10u10	147440-19 147441-19	ZFS	EFI	solaris	physical-block-size: 4096	6.1
10u8, 10u9, 10u10	Default	SVM Slice vXvm	SMI	solaris		6.1 6.0
10u8, 10u9, 10u10	Default	SVM Slice vXvm	EFI	solaris_efi		6.1 6.0
10u6, 10u7	Default	ZFS	EFI	solaris		6.0 5.1
10u6, 10u7	Default	SVM Slice vXvm	SMI	solaris		6.0 5.1
10u6, 10u7	Default	SVM Slice vXvm	EFI	solaris_efi		6.0 5.1
10u5	Default	ZFS	EFI	solaris_efi		6.0 5.1

OS Version	Kernel	Volume Manager	OS Label	LUN Type	(s)sd.conf changes	Host Utility
10u5	Default	SVM Slice vXvm	SMI	solaris		6.0 5.1
10u5	Default	SVM Slice vXvm	EFI	solaris_efi		6.0 5.1
10u4 and earlier	Default	ZFS	EFI	solaris_efi		5.1
10u4 and earlier	Default	SVM Slice vXvm	SMI	solaris		5.1
10u4 and earlier	Default	SVM Slice vXvm	EFI	solaris_efi		5.1

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 • <i>IBM System Storage N series Introduction and Planning Guide</i>
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 for 7-Mode</i> • <i>Data ONTAP SAN Administration Guide for 7-Mode</i> • <i>Data ONTAP Release Notes</i> • <i>Command Reference</i>
Verifying compatibility of a storage system with environmental requirements	<i>N series Introduction and Planning Guide</i>

If you need more information about...	Go to...
Upgrading Data ONTAP	<i>Data ONTAP Upgrade and Revert/Downgrade Guide for 7-Mode</i>
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, such as SVM, ZFS, or MPxIO	Refer to your operating system documentation. You can download Oracle manuals in PDF format from the Oracle website.
Working with Emulex	Refer to the Emulex documentation.
Working with QLogic	Refer to the QLogic documentation.
General product information, including support information	The IBM N series interoperability matrix website (accessed and navigated as described in Websites on page 9)

Related information

[Emulex partner site \(when this document was produced\)](#)

[QLogic partner site \(when this document was produced\)](#)

[Oracle documentation \(when this document was produced\)](#)

[Veritas Storage Foundation DocCentral](#)

Copyright and trademark information

Copyright ©1994 - 2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2013 IBM Corporation. All rights reserved.

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

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

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

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

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

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

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

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

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

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

Trademark information

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

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

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

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

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

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

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

Index

A

ALUA

- supported configurations [113](#)
- supported with MPxIO and FC [113](#)

APM

- available from Symantec [51](#)
- example of installing [58](#)
- example of uninstalling [56](#)
- installed with certain Veritas versions [51](#)
- installing [55](#)
- obtaining from Symantec [54](#)
- required in Veritas environments [52](#)
- uninstalling [56](#)
- uninstalling with pkgrm [52](#)
- used with Veritas [51](#)

array types

- displaying with sanlun [61](#)

ASL

- array type [60](#)
- available from Symantec [51](#)
- determining the ASL version [53](#)
- error messages [64](#)
- example of installing [58](#)
- example of uninstalling [56](#)
- installed with certain Veritas versions [51](#)
- installing [55](#)
- obtaining from Symantec [54](#)
- required in Veritas environments [52](#)
- uninstalling [56](#)
- uninstalling with pkgrm [52](#)
- used with Veritas [51](#)

ASUP

- collectinfo command [78](#)

B

brocade_info

- installed by Host Utilities [13](#)

C

cfgadm -al command [68](#)

CHAP

- configuring [38](#)
- secret value [38](#)

CHAP authentication [38](#)

cisco_info

- installed by Host Utilities [13](#)

collectinfo

- examples [79](#)
- format [78](#)
- options [78](#)
- overview [78](#)

configurations

- finding more information [135](#)
- supported for Host Utilities [119](#)

controller_info

- installed by Host Utilities [13](#)

D

device names

- generated using LUNs serial numbers [124](#)

direct-attached configurations

- supported by Host Utilities [119](#)

DMP

- working with iSCSI [37](#)

documentation

- finding more information [135](#)

drivers

- getting software [22](#)
- installing SANsurfer CLI [26](#)

dynamic discovery

- iSCSI SendTargets command [37](#)

E

EMLXemlxu

- FCA utilities [23](#)
- installing [24](#)

emlxs

- Solaris native drivers for Emulex HBAs [23](#)

Emulex

- changing HBA to SFS mode [83](#)
- downloading software and firmware [22](#)
- getting software [22](#)
- updating firmware and boot code [22](#)
- using Solaris native drivers [23](#)

error messages

- ASL [64](#)

example

installing ASL, APM *58*
 uninstalling ASL, APM *56*

F

fabric-attached configurations
 supported by Host Utilities *119*

fast recovery
 Veritas feature *61*

FC protocol
 ALUA and MPxIO *113*
 ALUA configurations *113*
 supported by Host Utilities *118*
 supported configurations *119*

fc_err_recov parameter
 required for Veritas *51*

FCA utilities
 EMLXemlxu *23*
 installing EMLXemlxu *24*
 managing Solaris native drivers *23*

FCode HBAs with native drivers
 updating *25*

Fibre Channel protocol *See* FC protocol

filer_info
 installed by Host Utilities *13*

finding more information *135*

firmware
 upgrades for native drivers *24*

format utility
 labeling LUNs *70*

G

GRUB installing information (X64) *104*

H

HBA
 displaying information with sanlun *76*

host settings
 software package *30*

Host Utilities
 ASL and APM for Veritas *51*
 ASL and APM requirements *52*
 contents *13*
 defined *13*
 downloading software packages *30*
 environments *14*
 finding more information *135*
 installation prerequisites *18*

installing *31*

iSCSI configuration *36*
 key setup steps *29*
 MPxIO environment *14*
 MPxIO stack overview *116*
 planning installation *18*
 software packages *29*
 support for non-English versions *115*
 supported configurations *119*
 uncompressing software packages *30*
 uninstalling Attach Kit 2.0 *35*
 uninstalling versions 5.x, 4.x, 3.x, 6.x *33, 34*
 upgrading *33*
 using Logical Domains *114*
 using VxVM *60*
 Veritas environment *14*
 Veritas stack overview *115*

host_config
 examples *42*
 options *40*

host_config command
 configuring all environment *40*

I

igroup
 uses WWPN *118*

igroup create command *67*

igroups
 creating *67*

information
 finding more *135*

installation
 downloading software packages *30*
 Host Utilities *31*
 iSCSI configuration *36*
 key setup steps *29*
 overview *18*
 planning *18*
 prerequisites *18*
 software packages *29*
 uncompressing software packages *30*

iSCSI protocol
 configuration *20, 36*
 configuring bidirectional CHAP *38*
 configuring discovery *37*
 discovering LUNs *68*
 implementing with Host Utilities *118*
 MPxIO *112*
 node names *36*

- recording node names [37](#)
- storage system IP address [37](#)
- supported configurations [119](#)
- troubleshooting [129](#)
- Veritas DMP [112](#)
- working with Veritas [37](#)

- ISNS discovery
- iSCSI [37](#)

L

- labeling LUNs [70](#)

- languages

 - support for non-English versions [115](#)

- Locally booted disks

 - copying data from (UFS) [96](#)

 - copying data from (ZFS) [98](#), [100](#)

- Logical Domains

 - using with Host Utilities [114](#)

- LPFC drivers

 - getting Emulex software [22](#)

- lun create command [67](#)

- lun map command [67](#)

- lun setup command

 - creating LUNs, igroups [67](#)

- LUNs

 - configuration overview [65](#)

 - configuring [20](#)

 - creating [67](#)

 - creating bootable [87](#)

 - creating SAN boot [81](#), [82](#)

 - creating Veritas boot LUN [81](#)

 - discovering when using iSCSI [68](#)

 - discovering with native drivers [69](#)

 - displaying with sanlun [73](#)

 - getting controller number (Sun native drivers) [68](#)

 - labeling [70](#)

 - labeling SAN boot [93](#)

 - LUN type and performance [67](#)

 - mapping [66](#)

 - partitioning SAN boot [93](#)

 - probing idle settings [63](#)

 - SAN bootable [102](#)

 - SAN booting on SPARC systems [107](#)

 - SAN booting on X64 systems [108](#)

 - types, OS labels and combinations [133](#)

M

- man pages

 - installed by Host Utilities [13](#)

- mcdata_info

 - installed by Host Utilities [13](#)

- mpathadm

 - verifying multipathing on MPxIO [124](#)

- MPIO

 - getting Host Utilities software package [30](#)

- MPxIO

 - ALUA [14](#)

 - ALUA configurations [113](#)

 - ALUA support [113](#)

 - checking multipathing [124](#)

 - environment for Host Utilities [14](#)

 - fast recovery with Veritas [61](#)

 - iSCSI protocol [112](#)

 - labeling LUNs [70](#)

 - multipathing [14](#)

 - preparing for Host Utilities [49](#)

 - protocols [14](#)

 - sd.conf variables [50](#)

 - setting up a stack [49](#)

 - ssd.conf variables [50](#)

 - stack overview [116](#)

 - troubleshooting [131](#)

 - volume managers [112](#)

- multipathing

 - options [112](#)

 - using sanlun with Veritas [61](#)

 - verifying on MPxIO [124](#)

N

- network-attached configurations

 - supported by Host Utilities [119](#)

- node names

 - iSCSI [36](#)

 - recording [37](#)

- non-English versions supported [115](#)

- nSANity

 - installing [131](#)

O

- OpenBoot [82](#)

- outpassword

 - CHAP secret value [38](#)

P

- paths

- displaying with sanlun [75](#)
- performance
 - affected by LUN type [67](#)
- polling interval
 - recommended values [62](#)
- problems
 - checking troubleshooting [120](#)
- publications
 - finding more information [135](#)

Q

- QLogic
 - creating FCode compatibility [85](#)
 - downloading and extracting software [25](#)
 - getting HBA software [22](#)
 - qlc [25](#)
 - SANsurfer CLI [26](#)
- qlogic_info
 - installed by Host Utilities [13](#)

R

- requirements
 - finding more information [135](#)
- restore policy
 - recommended value [62](#)
- Root encapsulation
 - enabling [110](#)

S

- SAN booting
 - advantages [114](#)
 - changing Emulex HBA to SFS mode [83](#)
 - configuration steps [82](#)
 - configuring host to boot on X64 systems [108](#)
 - configuring hosts [107](#)
 - configuring hosts on SPARC based systems [107](#)
 - creating LUNs [81](#), [82](#)
 - creating Veritas boot LUNs [81](#)
 - FCodes compatibility with QLogic HBA [85](#)
 - gathering information for Native MPxIO [90](#)
 - gathering information for Veritas DMP [87](#)
 - gathering source disk information for [92](#)
 - making the LUN bootable [102](#)
 - partitioning LUNs [93](#)
 - setting up Oracle native HBA for [83](#)
- SAN Toolkit
 - getting Host Utilities software package [30](#)

- san_version command
 - installed by Host Utilities [13](#)
- sanlun utility
 - displaying array types [61](#)
 - displaying HBA information [76](#)
 - displaying LUNs [73](#)
 - displaying multipathing for Veritas [61](#)
 - displaying paths [75](#)
 - installed by Host Utilities [13](#)
 - output [75](#)
 - verifying multipathing on MPxIO [124](#)
- SANsurfer CLI
 - installing [26](#)
- sd.conf
 - recommended values for MPxIO [50](#)
- software packages
 - downloading Host Utilities software [30](#)
 - installing Host Utilities [31](#)
 - SPARC processor [29](#)
 - uncompressing [30](#)
 - x86/64 processor [29](#)
- Solaris
 - EFI labeling scheme [67](#)
 - Host Utilities [13](#)
 - Host Utilities environments [14](#)
 - labeling LUNs [70](#)
 - support for non-English language versions [115](#)
 - using Logical Domains [114](#)
- Solaris HBAs
 - working with [22](#)
- Solaris Host Utilities [13](#)
- Solaris native drivers
 - FCA utilities [23](#)
 - getting Emulex HBAs [22](#)
- solaris_efi
 - LUN type [67](#)
- solaris_info
 - installed by Host Utilities [13](#)
- SPARC
 - configuring host to boot from a SAN boot LUN [107](#)
 - installing the boot block [102](#)
 - OpenBoot [82](#)
- SPARC processor
 - software package [29](#)
- ssd_config.pl script
 - installed by Host Utilities [13](#)
- ssd.conf file
 - recommended values for MPxIO [50](#)
 - recommended values with Veritas DMP [48](#)
- static discovery

- iSCSI [37](#)
- storage system
 - iSCSI discovery [37](#)
 - outpassword and CHAP [38](#)
- storage systems
 - displaying path information with sanlun [75](#)
- Sun native drivers
 - discovering LUNs [69](#)
 - getting HBA controller number [68](#)
- Sun Volume Manager *See* SVM
- SVM
 - managing LUNs [72](#)
 - MPxIO [112](#)
- Symantec
 - getting the ASL and APM [54](#)
 - provides ASL, APM [51](#)

T

- troubleshooting
 - finding information [120](#)
 - finding iSCSI information [129](#)
 - finding more information [135](#)
 - possible MPxIO issues [131](#)

U

- UFS
 - copying data from locally booted disk [96](#)
- uninstalling
 - Attach Kit 2.0 [35](#)
 - Host Utilities versions 5.x, 4.x, 3.x, 6.x [33](#), [34](#)
- upgrading
 - host utilities [33](#)

V

- Veritas
 - ALUA configurations [113](#)
 - APM [51](#)
 - ASL [51](#)
 - ASL and APM requirements [52](#)
 - configuring iSCSI [37](#)
 - creating a SAN boot LUN [81](#)
 - displaying multipathing [61](#)
 - drivers [14](#)
 - enabling root encapsulation [110](#)

- environment for Host Utilities [14](#)
- fast recovery [61](#)
- fc_err_recov parameter [51](#)
- getting driver software [22](#)
- iSCSI protocol [112](#)
- labeling LUNs [70](#)
- multipathing [14](#)
- path age settings [63](#)
- preparing for Host Utilities [47](#)
- protocols [14](#)
- restore daemon settings [62](#)
- setup issues [14](#)
- Solaris native drivers [23](#)
- ssd.conf variables [48](#)
- stack overview [115](#)
- using VxVM to display LUN paths [60](#)
- volume manager [112](#)
- volume management
 - managing LUNs [72](#)
- volume managers
 - MPxIO [112](#)
 - Veritas [112](#)
- VxVM
 - displaying LUN paths [60](#)
 - managing LUNs [72](#)
 - MPxIO [112](#)
 - Veritas volume manager [112](#)

W

- WWPN
 - getting with sanlun fcp show adapter command [118](#)
 - required for igroups [118](#)

X

- X64
 - installing GRUB information [104](#)
- x86/64 processor
 - software package [29](#)

Z

- ZFS
 - copying data from locally booted disk [98](#), [100](#)
 - managing LUNs [72](#)
 - volume manager with MPxIO [112](#)



NA 210-05946_A0, Printed in USA

GC52-1348-02

