

FCP IBM® AIX® Host Utilities 4.0 (VERITAS Storage Foundation) **Installation and Setup Guide**

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 215-02019_A0
September 2006

Copyright and trademark information

Copyright information

Copyright © 1994–2006 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

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 Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL NETWORK APPLIANCE 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.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. 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 Network Appliance.

The product described in this manual may be protected by one or more U.S. 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

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

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

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	vii
Chapter 1	Installation and Setup Overview	1
	FCP IBM AIX Host Utility installation checklist	2
	Completing the pre-installation tasks.	6
	Where to go for more information	11
Chapter 2	Installing the Software Package	13
	Obtaining the software packages	14
	Installing the software packages	18
	Uninstalling the software packages.	29
Chapter 3	Installing the ASL and APM	35
	Overview of the Array Support Library and Array Policy Module	36
	Installing the ASL and APM	37
Chapter 4	Migrating a Storage Cluster to Single Image Mode.	41
	Deciding whether to migrate to single image mode	42
	Understanding cfmode	43
	Tasks to complete before you begin	47
	Upgrading to single image mode	49
Chapter 5	Creating, Configuring, and Managing LUNs	55
	Preparing the storage system	56
	Creating and mapping igroups and LUNs	57
	LUN configuration overview.	62
	Configuring LUNs with native AIX LVM.	64
	Configuring LUNs with VxVM and DMP multipathing software.	77
	Displaying LUN and HBA information	83
	Using sanlun to display host LUN information	84

Displaying host HBA information 87

Chapter 6

Troubleshooting 89

Troubleshooting checklists 90

Installation fails because system does not meet requirements 92

Commands you can use to troubleshoot problems 93

SAN Toolkit diagnostic utilities 94

Change the AIX HBA's attributes 100

Change the HBA queue depth with no LUNs mapped 104

Change the value for the queue depths on the disk. 105

General troubleshooting 106

Index 107

Preface

About this guide	This guide describes how to integrate an IBM® AIX® host with a Network Appliance™ storage system with Data ONTAP™ software and the Fibre Channel Protocol (FCP) for SCSI.
Audience	This guide is for system installers and administrators who are familiar with the AIX operating system. It assumes that you understand how to configure and administer a storage system and use FCP, to store and retrieve data on the system.
Terminology	This guide uses the following terms: ◆ <i>igroup</i> refers to a collection of unique identifiers, either FCP WWPNs in a SCSI network or iSCSI node names of initiators (hosts) in an IP network, that are given access to LUNs when they are mapped to those LUNs. ◆ <i>LUN</i> (Logical Unit Number) refers to a logical unit of storage identified by a number. ◆ <i>LUN ID</i> refers to the numerical identifier for a LUN. ◆ <i>Initiator</i> refers to the system component that originates an I/O command over an I/O bus or network. ◆ <i>Target</i> refers to the storage system that receives an FCP I/O request and serves the data requested by the initiator. ◆ <i>Type</i> refers to pressing one or more keys on the keyboard. ◆ <i>Enter</i> refers to pressing one or more keys and then pressing the Enter key.
Command conventions	You can enter storage system commands on either the system console or from any host system that can access the system through a Telnet session.
Keyboard conventions	When describing key combinations, this guide uses the hyphen (-) to separate individual keys. For example, “Ctrl-d” means pressing the “Control” and “d” keys simultaneously. Also, this guide uses the term “Enter” to refer to the key that generates a carriage return, although the key is named “Return” on some keyboards.

Typographic conventions

The following table describes typographic conventions used in this guide.

Convention	Type of information
<i>Italic font</i>	Words or characters that require special attention. Placeholders for information you must supply. For example, if the guide says to enter the <code>arp -d hostname</code> command, you enter the characters <code>arp -d</code> followed by the actual name of the host. Book titles in cross-references.
<code>monospaced font</code>	Command and daemon names. Information displayed on the system console or other computer monitors. The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless you must type it in upper case.

Special messages

This guide contains the following types of special messages:

Note

A note contains important information that helps you install or operate the system efficiently.

Caution

A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

WARNING

A warning contains instructions that you must follow to avoid personal injury.

About this chapter

This chapter provides a brief overview of the installation process and the tasks you need to complete before you install the Host Utility.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“FCP IBM AIX Host Utility installation checklist”](#) on page 2
- ◆ [“Completing the pre-installation tasks”](#) on page 6
- ◆ [“Where to go for more information”](#) on page 11

FCP IBM AIX Host Utility installation checklist

List of tasks to perform

Installing the FCP IBM AIX Host Utility and setting up your system involves a number of steps that are performed on both the storage system and the host. The following is a high-level checklist of the tasks you must perform and the machines involved.

Step	Action	Machine	Details
Preparing to install the software package			
1	Review the information on the supported components and plan your configuration.		<i>Compatibility and Configuration Guide for NetApp FCP and ISCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml .
2	Verify that your configuration meets the system requirements and you have the correct components. Note _____The Host Utilities package does not include the host HBA.	Host	“ Verify that the host meets system requirements ” on page 6
3	Confirm that you have the correct operating system version and filesets.	Host	“ Check your operating system version ” on page 7
4	Verify that the storage system is set up and configured.	Storage System	Refer to: ◆ Installation and setup instructions that ship with your hardware ◆ <i>Fibre Channel Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FC_Configuration_Guide.pdf ◆ “Verifying that the storage system is configured” on page 8

Step	Action	Machine	Details
5	Verify that the switch, if you're using a fabric connection, is set up correctly and, if appropriate for your system, zoned.	Switch, Storage System, Host	Refer to: ◆ Installation and setup instructions that ship with your hardware ◆ <i>Fibre Channel Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FC_Configuration_Guide.pdf ◆ <i>Switch Compatibility Matrix</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_switch.shtml.
6	Verify that the power to the switches, disk shelves, storage systems, and hosts is turned on in the correct order.	Switch, Disk Shelves, Storage System, Host	“Verifying that the switch is installed and configured” on page 9
7	Verify that the host and storage system can communicate.	Storage System, Host	“Verifying the host and storage system can communicate” on page 10
Installing the software package			
8	Get a copy of the Host Utilities host settings software package and the SAN Toolkit software package.	Host	“Ways to obtain the software packages” on page 14
9	If you downloaded the software packages, uncompress them and extract the files. If you are using the FCP IBM AIX Host Utilities CD-ROM, mount the CD-ROM.	Host	Refer to: ◆ “Downloading the software” on page 15 ◆ “Getting the software from the CD-ROM” on page 16

Step	Action	Machine	Details
10	Install the host settings software package and the SAN Toolkit. If your configuration requires multipathing, install the following: ◆ SAN Toolkit ◆ VERITAS Storage Foundation 4.0 — available from Symantec. This solution requires a Symantec certified Array Support Library (ASL) and Array Policy Module (APM), which you obtain from the Symantec Web site. Note_____ You must install VERITAS Volume Manager (VxVM) before you install the ASL and APM software.	Host	Refer to: ◆ “Installing the software packages” on page 18 ◆ “Configuring LUNs with native AIX LVM” on page 64 ◆ “Configuring LUNs with VxVM and DMP multipathing software” on page 77. Also, there is a direct link to ASL and APM on the Symantec Web site from the FCP Host Compatibility Matrix. ◆ “Installing the ASL and APM” on page 35
11	Make sure the HBA initiator queue depth is appropriate for your system.	Host	“Verifying the HBA initiator queue depth” on page 28
Creating and configuring LUNs			
12	(Recommended) If you have a storage cluster, migrate your cfmode setting to single image mode. If you have a direct-attached storage, verify that the media type is correct.	Storage System	Refer to: ◆ Chapter 4, “Migrating a Storage Cluster to Single Image Mode,” on page 41 ◆ “Setting the media type” on page 56
13	Create igroup(s) using AIX as the value for the ostype attribute.	Host, Storage System	“Creating an igroup” on page 57
14	Create LUN(s) that map to the igroup.	Storage System	“Creating and mapping LUNs” on page 59

Step	Action	Machine	Details
15	Use <code>cfgmgr</code> to get the host to discover the new LUNs.	Host	“ Getting the host to discover the new LUNs ” on page 64
16	Configure the LUNs ◆ If your configuration does not require multipathing, use native AIX LVM to configure LUNs. ◆ If your configuration requires multipathing, use VERITAS software to configure LUNs.	Host	Refer to: ◆ “Configuring LUNs with native AIX LVM” on page 64 ◆ “Configuring LUNs with VxVM and DMP multipathing software” on page 77.
17	Display information about the LUNs and HBA.	Host	Refer to: ◆ “Using sanlun to display host LUN information” on page 84 ◆ “Displaying host HBA information” on page 87

Completing the pre-installation tasks

Before you install the Host Utilities

There are several tasks you should perform before you install the Host Utilities software. The tasks ensure that you meet the requirements for the Host Utilities and have your system set up correctly.

Verify that the host meets system requirements

Check your host setup against the system requirements. Support for new components is added on an on-going basis. For a current and complete list of new and existing HBAs, platforms, applications, and drivers, see the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Use Dynamic LPAR, if appropriate

Dynamic Logical Partitioning (DLPAR) is an IBM AIX 5.2 and 5.3 feature on IBM's pSeries servers. With DLPAR, you can move hardware resources such as processors, memory, and I/O slots from one logical partition to another without disruption. Refer to IBM AIX documentation for details about DLPAR.

Check your operating system version

If you're not sure you have a supported version of the operating system, enter the following command:

```
oslevel -s
```

Result: Information about your operating system and maintenance or technology level appears. This command gives you basic information about the operating system level, but it may not be complete. You may need additional filesets.

Example: This command displays information on your operating system level.

```
# oslevel -s
5300-04-03
```

Check the version of the prerequisite filesets

Use the `lslpp` command to check the filesets. The command format is:

```
lslpp -L <fileset_name>
```

`-L` displays information about filesets installed on the host.

`fileset_name` displays information about that fileset on the host.

Result: The output from this command tells you which filesets are installed on the host or provides information about a fileset on the host.

Example: The following example displays information about the fileset `devices.pci.df1000f7.com`.

```
# lslpp -L devices.pci.df1000f7.com
  Fileset                                Level  State  Type  Description
(Uninstaller)
-----
devices.pci.df1000f7.com 5.2.0.40 C      F      Common PCI FC Adapter
Device Software
State codes:
A -- Applied.
B -- Broken.
C -- Committed.
E -- EFIX Locked.
O -- Obsolete. (partially migrated to newer version)
? -- Inconsistent State...Run lppchk -v.
Type codes:
F -- Installp Fileset
P -- Product
C -- Component
T -- Feature
```

Verifying the correct host HBA is installed

Supported HBAs should be installed before you install the software package. Complete the following steps to display information about HBAs on your system and, if necessary, to install an HBA.

Step	Action
1	To display HBA information, enter: <pre>lsdev -Cc adapter</pre> Results: The system displays information similar to the following. In this case, fcs0 and fcs1 are fibre channel adapters. <pre># lsdev -Cc adapter ent0 Available 03-08 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902) ent1 Available 03-09 2-Port 10/100/1000 Base-TX PCI-X Adapter (14108902) fcs0 Available 05-08 FC Adapter fcs1 Available 09-08 FC Adapter ide0 Available 07-08 ATA/IDE Controller Device sisscsia0 Defined 04-08 PCI-X Dual Channel Ultra320 SCSI Adapter sisscsia1 Available 08-08 PCI-X Dual Channel Ultra320 SCSI Adapter usbhcc0 Available 02-08 USB Host Controller (33103500) usbhcc1 Available 02-09 USB Host Controller (33103500) vsa0 Available LPAR Virtual Serial Adapter vsa1 Available LPAR Virtual Serial Adapter</pre>
2	If you need to install an HBA, refer to the documentation that came with your HBA.

Verifying that the storage system is configured

Before you configure initiator groups and LUNs on the storage system, make sure you have completed the following:

1. Set up the AIX host and the initiator HBAs.
2. Configured the storage system and target HBAs (if needed).
3. Cabled the storage to the fabric according to your configuration topology.
4. Licensed and enabled the FCP protocol on the storage system.
5. Created working volumes and qtrees (if desired) on the storage system.

Note

It is recommended that you create volumes and LUNs with space reservations enabled.

Refer to “*Where to go for more information*” on page 11 for a list of useful documents and web sites.

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 Installation and Setup instructions that shipped with your hardware. You should have completed the following tasks:

- ◆ Installed the switch in a system cabinet or rack.
- ◆ Checked that the switch firmware is one of the versions supported with this Host Utility. See the *Fibre Channel Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPCConfigurationGuide.pdf.
- ◆ Turned on power to the switch.
- ◆ Configured the network parameters for the switch, including its serial parameters and IP address.

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 results in the following:

- ◆ Configuring the host is simplified.
- ◆ Information becomes more manageable. The output from host tools such as `iostat` and `lsdev` 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 *Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Correct order for powering up the equipment

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. Throughout the procedures that follow, this is the order you must use when powering on the equipment:

- ◆ Configured Fibre Channel switches
It can take several minutes for the switches to boot.
- ◆ Disk shelves
- ◆ Storage systems
- ◆ Host

Verifying the host and storage system can communicate

To verify that the host can issue commands to the storage system through the storage system's console or command line, complete the following step.

Step	Action
1	Issue a command from the host to the storage system. One way to do this is to use a remote login, such as <code>telnet</code>, to connect with the storage system. Once you have logged in, enter a command at the filer prompt.

Where to go for more information

Host and storage system information

The documents listed below contain information about host and storage system requirements, supported configurations, best practices, installation procedures and troubleshooting.

If you need information about...	See the following product documentation...
The latest supported configurations	◆ <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Configuring the SAN environment	◆ <i>Network Appliance Fibre Channel Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPCConfigurationGuide.pdf
Provisioning and space requirements Managing LUNs Managing initiator groups Using Data ONTAP	◆ <i>Data ONTAP™ Block Access Management Guide</i> ◆ <i>Data ONTAP Release Notes</i> ◆ <i>Data ONTAP Upgrade Guide</i> Note— You can download PDFs of these documents from the “Access and Order Product Documentation” page on the NOW site at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Verifying compatibility of storage systems with environmental requirements	◆ <i>Site Requirements Guide</i>, which available online at http://now.netapp.com/NOW/knowledge/docs/docs.shtml

About this chapter

This chapter describes procedures for installing the FCP IBM AIX Host Utility and optional multipathing software.

Note

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

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Obtaining the software packages”](#) on page 14
- ◆ [“Installing the software packages”](#) on page 18
- ◆ [“Uninstalling the software packages”](#) on page 29

Obtaining the software packages

Software packages

The FCP IBM AIX Host Utility includes:

- ◆ The FCP IBM AIX host settings software package — contains NetApp's FCP device definitions that are required for FCP IBM AIX configurations.
- ◆ The NetApp SAN Toolkit software package — contains the `sanlun` utility and diagnostic scripts. The `sanlun` utility displays configuration information about LUNs and HBAs. The diagnostic scripts collect information about your system if a problem occurs. Customer support may ask you to run the `sanlun` utility or a diagnostic script if a problem occurs.

You must download both the host settings software package and the SAN Toolkit software package.

Configurations that include multipathing require Array Support Library 4.0 (ASL) and Array Policy Module (APM) for IBM AIX, which you obtain from Symantec.

Ways to obtain the software packages

You can obtain the software packages in two ways:

- ◆ Download them from the NOW site. The FCP IBM AIX Host Utility on the NOW site provides the host settings software package and the SAN Toolkit software package. You must download both packages. For information about downloading these files, see “[Downloading the software](#)” on page 15.
- ◆ Get them from the FCP IBM AIX Host Utility CD-ROM. Information on getting the software packages from the CD-ROM is detailed in the section “[Getting the software from the CD-ROM](#)” on page 16.

Note

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

Downloading the software

To download the software from the NOW site, complete the following steps.

Step	Action
1	Log in to the NOW site at http://now.netapp.com . In the fields on the right side of the page, complete the following steps: 1. Enter your user name and password. 2. From the Select Start Page list box, click Download Software. 3. Click the Login button.
2	On the Download Software page, go to the SAN Host Utility (FCP) product row of the table and select AIX from the Select Platform drop-down list.
3	Click Go!
4	Follow the prompts to reach the Software Download page.
5	Download the NetApp host settings software file and the SAN Toolkit software file to a local directory on your host. You must download both files. Note _____ If you download the files to a machine other than the AIX host, you must move them to the host. _____

Uncompressing the software file you downloaded

To uncompress the software files you downloaded from the NOW site, complete the following steps.

Step	Action
1	Change to the directory where you saved the downloaded software and enter the following commands to uncompress the FCP IBM AIX host settings package and the SAN Toolkit package: <pre># zcat _aix_fcp_4.0.tar.Z tar -xvf - # zcat santoolkit_aix_ibm_v3.0.0.tar.Z tar xvf -</pre> Result: The extracted software is placed in the directory where you downloaded the software. Example: Executing the <code>zcat</code> and <code>tar</code> commands with the host settings file and the SAN Toolkit file produces output similar to the following. <pre>HostA# zcat NetApp_aix_fcp_4.0.tar.Z tar xvf - x ./NetApp_aix_fcp_4.0/NetApp.fcp_attach_kit, 22528 bytes, 44 media blocks. x ./NetApp_aix_fcp_4.0/NetApp.fcp_toolkit, 1005568 bytes, 1964 media blocks. x ./NetApp_aix_fcp_4.0/install, 37326 bytes, 73 media blocks. x ./NetApp_aix_fcp_4.0/uninstall, 17844 bytes, 35 media blocks. HostA# zcat santoolkit_aixv3.0.0.tar.Z tar xvf - x ./NetApp.SAN_toolkit, 1146880 bytes, 2240 media blocks. HostA#</pre>
2	Go to “ Installing the software packages ” on page 18.

Getting the software from the CD-ROM

To get the software from the CD-ROM, complete the following steps.

Step	Action
1	Insert the CD-ROM into the CD-ROM drive.
2	Change to the root directory of the CD-ROM: <pre># cd <CDROM_mountpoint></pre>

Step	Action
3	Copy the host settings software package and the SAN Toolkit package from the CD-ROM's root directory to a temporary location on the host. Example: If you mount the CD-ROM on /mnt and want to place the SAN Toolkit file in the /tmp directory, you might enter commands similar to the following: <pre># mount -v cdrfs -o ro /dev/cd0 /mnt # cp /mnt/santoolkit_aix_v3.0.0.tar.Z /tmp # cd /tmp</pre>
4	Go to “ Installing the software packages ” on page 18.

Installing the software packages

Tasks involved in installing the Host Utility software

You can install the host settings software and the SAN Toolkit software using either the AIX System Management Interface Tool (SMIT) or the `install -p` command.

Note

This section uses SMIT as its installation example.

Upgrading from an earlier attach kit

If you have an FCP AIX Host Attach Kit or a previous version of the Host Utilities installed, you should remove it before you install a new version of the Host Utilities. See “[Uninstalling the software packages](#)” on page 29.

Note

Previous versions of the Host Utility were called Attach Kits.

Using SMIT to
install the software

You must install both the host setting software and the SAN Toolkit software. To do this, complete the following steps.

Note_____

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

Step	Action
1	Make sure you are logged in as root.
2	Start SMIT by entering the following command: # smit install
3	At the screen that appears, select the option titled “Install and Update Software.” Example: When you start SMIT, it displays the following screen. On this screen, the “Install and Update Software” option is the first menu option. <pre># smitty install +-----+ Software Installation and Maintenance Move cursor to desired item and press Enter. Install and Update Software List Software and Related Information Software Maintenance and Utilities Software Service Management Network Installation Management EZ NIM (Easy NIM Tool) System Backup Manager Alternate Disk Installation EFIX Management F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
4	At the next screen, select the Install Software menu option. Example: The following is an example of the Install and Update Software screen. <pre> +-----+ Install and Update Software Move cursor to desired item and press Enter. Install Software Update Installed Software to Latest Level (Update All) Install Software Bundle Update Software by Fix (APAR) Install and Update from ALL Available Software F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
5	At the Install Software screen, specify the location of the software in one of the following ways: ◆ Manually enter the location by providing the following information. ❖ If you're installing from the CD-ROM, enter the CD-ROM drive. ❖ If you're installing from the host machine, enter the path to the software package (for example, /var/tmp/ntap_aix_veritas_4.0.0/NetApp.veritas_attach_kit). ◆ Press F4 to display a list of options. If you want to use the F4 method, complete the following steps: 1. Press F4. 2. At the prompt for software to install, enter NetApp.veritas_attach_kit. 3. At the prompt asking whether you want to continue or cancel, press Enter to complete the installation. Example 1: The following is an example of entering the path to the software package when you're at the Install Software screen. <pre> +-----+ Install Software Type or select a value for the entry field. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software [/var/tmp/ntap_aix_veritas_4.0.0/NetApp.veritas_attach_kit]+ F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
	Example 2: After you enter path to the software package, SMIT displays the following screen. This is the screen where you enter name of the software package, Netapp.veritas_attach_kit, and press Enter. <pre> +-----+ Install Software Type or select values in entry fields. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software [/var/tmp/ntap_aix_veritas_4.0.0/NetApp.veritas_attach_kit] + * SOFTWARE to install [NetApp.veritas_attach_kit]+ PREVIEW only? (install operation will NOT occur) no + COMMIT software updates? yes + SAVE replaced files? no + AUTOMATICALLY install requisite software? yes + EXTEND file systems if space needed? yes + OVERWRITE same or newer versions? no + VERIFY install and check file sizes? no + Include corresponding LANGUAGE filesets? yes + DETAILED output? no + Process multiple volumes? yes + ACCEPT new license agreements? no + Preview new LICENSE agreements? no + F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
6	SMIT displays a warning, asking you to confirm that you want to install the software. Press Enter to continue. Example 1: The following is an example of entering the path to the software package when you're at the Install Software screen. <pre>x*****x x ARE YOU SURE? x x x x Continuing may delete information you may want x x to keep. This is your last chance to stop x x before continuing. x x Press Enter to continue. x x Press Cancel to return to the application. x x x x Flx F1=Help F2=Refresh F3=Cancel x x Esx Esc+8=Image Esc+0=Exit Enter=Do x x*****x</pre>

Step	Action
7	SMIT displays information about the installation. Note_____ By default, the host settings software installs the filesets for both FCP and iSCSI. This is not a problem. If you are using only FCP, ignore the iSCSI fileset. _____ Example of a successful installation: The following is an example of output you might see when an installation successfully completes. <pre> COMMAND STATUS Command: running stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.veritas_attach_kit.pcm0dm 4.0.0.0 NetApp.veritas_attach_kit.iscsi 4.0.0.0 NetAppveritas_attach_kit.fcp 4.0.0.0 NetApp.veritas_attach_kit.config 4.0.0.0 << Copyright notice for NetApp.veritas_attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2006. All rights reserved. << End of copyright notice for NetApp.veritas_attach_kit >>. . . . Finished processing all filesets. (Total time: 5 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23825 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. * * * A T T E N T I O N * * * System boot image has been updated. You should reboot the system as soon as possible to properly integrate the changes and to avoid disruption of current functionality.</pre>

Step	Action
	<pre> installp: bosboot process completed. +-----+ Summaries: +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.veritas_attach_kit.pcm 4.0.0.0 USR APPLY SUCCESS NetApp.veritas_attach_kit.iscsi 4.0.0.0 USR APPLY SUCCESS NetApp.veritas_attach_kit.fcp 4.0.0.0 USR APPLY SUCCESS NetApp.veritas_attach_kit.conf 4.0.0.0 USR APPLY SUCCESS COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.veritas_attach_kit.config 4.0.0.0 << Copyright notice for NetApp.veritas_attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2006. All rights reserved. << End of copyright notice for NetApp.veritas_attach_kit >>. . . . Finished processing all filesets. (Total time: 5 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23825 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. installp: bosboot process completed. +-----+ Summaries: +-----+ Installation Summary ----- </pre>

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

NetApp.SAN_toolkit.scripts	3.0.0.0	USR	APPLY	SUCCESS
NetApp.SAN_toolkit.sanlun	3.0.0.0	USR	APPLY	SUCCESS
#				

Verifying the HBA initiator queue depth

It’s a good practice to verify the value of the HBA initiator queue depth. To check the queue depth, perform the following step.

Step	Action
1	Check the HBA queue depth by using the <code>lsattr -El</code> command. Example: For HBA <code>fcs0</code>, you would enter the following command: <pre># lsattr -El fcs0</pre> Note _____ It is recommended to use a value of 48 for the HBA queue depth. _____

Uninstalling the software packages

Before you uninstall previous versions of software

- Before you uninstall a previous version of the software, make sure you do the following;
- ◆ Unmount the file systems that contain the volume groups mapped to storage system LUNs.
 - ◆ Quiesce and export the volume groups.
 - ◆ Stop multipathing software.
 - ◆ Remove the existing NetApp LUN handles.
- These procedures disconnect NetApp LUNs from the host before you uninstall the software. For details, see “[Upgrading to single image mode](#)” on page 49.

Uninstalling the software from an AIX host

You can use SMIT to uninstall the Host Utilities software.

Note_____

You should always uninstall the current Host Utilities software before installing a new version.

Uninstalling the software using SMIT

To uninstall the software using SMIT, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Start SMIT by entering the following command: # smit

Step	Action
3	Select the “Software Installation and Maintenance” menu option. Example: When you start SMIT, it displays the following screen. <pre> +-----+ root> # smit System Management Move cursor to desired item and press Enter. Software Installation and Maintenance Software License Management Devices System Storage Management (Physical & Logical Storage) Security & Users Communications Applications and Services Print Spooling Problem Determination Performance & Resource Scheduling System Environments Processes & Subsystems Applications Cluster System Management Using SMIT (information only) F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
4	At the screen that appears, select the Software Maintenance and Utilities menu option. Example: The following is an example of the Software Installation and Maintenance screen. <pre> +-----+ Software Installation and Maintenance Move cursor to desired item and press Enter. Install and Update Software List Software and Related Information Software Maintenance and Utilities Network Installation Management System Backup Manager F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
5	At the next screen, select the Remove Installed Software menu option. Example: The following is an example of the Software Maintenance and Utilities screen. <pre> +-----+ Software Maintenance and Utilities Move cursor to desired item and press Enter. Commit Applied Software Updates (Remove Saved Files) Reject Applied Software Updates (Use Previous Version) Remove Installed Software Copy Software to Hard Disk for Future Installation Check Software File Sizes After Installation Verify Software Installation and Requisites Clean Up After Failed or Interrupted Installation F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
6	Remove the software in one of the following ways: ◆ Enter the package name (NetApp.veritas_attach_kit). (Make sure that the “Preview only” option is set to no.) ◆ Press F4 to display a list of names. If you use F4, complete the following steps: 1. Press F4. 2. Scroll down the list of names until you reach NetApp.veritas_attach_kit. 3. Select NetApp.veritas_attach_kit and press Enter. 4. At the prompt asking whether you want to continue or cancel, press Enter to complete the uninstall. Example 1: The following is an example of the Remove Installed Software screen. Note_____ By default PREVIEW only? is set to yes. You must change it to no if you want to uninstall. <pre> +-----+ Remove Installed Software Type or select values in entry fields. Press Enter AFTER making all desired changes. [Entry Fields] * SOFTWARE name [NetApp.veritas_attach_kit.config NetApp.veritas_attach_kit.fcp NetApp.veritas_attach_NetApp.veritas_attach_kit-> + PREVIEW only? (remove operation will NOT occur) no + REMOVE dependent software? no + EXTEND file systems if space needed? no + DETAILED output? no + F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
	Example of successful uninstall: The following output appears when you successfully uninstall the software. <pre> Command: running stdout: yes stderr: no Before command completion, additional instructions may appear below. 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. +-----+ +-----Deinstalling Software-----+ +-----+ installp: DEINSTALLING software for: NetApp.veritas_attach_kit.config 4.0.0.0 NetApp.veritas_attach_kit.fcp 4.0.0.0 NetApp.veritas_attach_kit.iscsi 4.0.0.0 NetApp.veritas_attach_kit.pcmo dm 4.0.0.0 Finished processing all filesets. (Total time: 4 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23818 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. * * * A T T E N T I O N * * * System boot image has been updated. You should reboot the system as soon as possible to properly integrate the changes and to avoid disruption of current functionality. installp: bosboot process completed. +-----+ +-----Summaries-----+ +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.veritas_attach_kit.conf i 4.0.0.0 USR DEINSTALL SUCCESS NetApp.veritas_attach_kit.fcp 4.0.0.0 USR DEINSTALL SUCCESS NetApp.veritas_attach_kit.isc s i 4.0.0.0 USR DEINSTALL SUCCESS NetApp.veritas_attach_kit.pcmod 4.0.0.0 USR DEINSTALL SUCCESS </pre>

About this chapter

If you are using VERITAS Storage Foundation 4.0 for multipathing, you must install and configure the Symantec Array Support Library (ASL) and Array Policy Module (APM). This chapter describes how to install and configure the ASL and APM and how these components affect multipathing and failover.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Overview of the Array Support Library and Array Policy Module”](#) on page 36
- ◆ [“Installing the ASL and APM”](#) on page 37

Overview of the Array Support Library and Array Policy Module

What the Array Support Library is

The Array Support Library (ASL) is an ONTAP-qualified library that provides information about storage array attributes and multipathing configuration to the Device Discovery Layer (DDL) and Dynamic Multipathing (DMP) components of VERITAS Volume Manager™ (VxVM).

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

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

Note

You cannot use storage systems simultaneously as JBODs and vendor arrays. If you install the ASL, storage systems cannot be configured in VxVM as JBODs. They are reported as storage arrays unless you explicitly exclude them by using the `vxddladm exclude array` command.

What the Array Policy Module is

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

Installing the ASL and APM

Obtaining the ASL and APM

The ASL and APM are not distributed with the attach kit. You obtain the ASL and APM from the Symantec Web site. The APM is distributed as a point patch to Volume Manager 4.0.3.3 (Point-Patch_3).

For a direct link to the ASL and APM on the Symantec Web site, see the FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

Note
Before you begin, make sure you obtain the ASL TechNote, which you can review at the Symantec Web site. The TechNote contains Symantec’s instructions for installing the ASL

Installing the ASL and APM

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

Step	Action	
1	Make sure you are logged into the host as the root user.	
2	If you have VxVM installed, stop all applications using VxVM volumes. Make sure that there are no VxVM volumes open, mounted, or in use.	
3	If you...	Then...
	You have VxVM 4.0 and patch MP3	Proceed to Step 4 .
	Do not have VxVM 4.0 and patch MP3	Install VxVM 4.0 and MP3 according to the procedures in the VxVM documentation and MP3 README file. Proceed to Step 4 .

Step	Action
4	For each host HBA (<i>fscsin</i>), enable the AIX FastFail feature. Enabling the FastFail feature improves I/O failover handling in the NetApp environment. Example input: <pre>chdev -l fscsi0 -a fc_err_recov=fast_fail -P chdev -l fscsi1 -a fc_err_recov=fast_fail -P</pre>
5	For each host HBA (<i>fscsin</i>), enable Dynamic Tracking. Example input: <pre>chdev -l fscsi0 -a dyntrk=yes -P chdev -l fscsi1 -a dyntrk=yes -P</pre>
6	Install the ASL according to the directions provided in the ASL tech note on the Symantec Web site.
7	Install the Volume Manager 4.0.3.3 (Point-Patch_3) point patch, which contains the APM. Use the procedures documented in the README file that ships with the point patch.
8	Reboot the host.
9	Verify that FastFail and Dynamic Tracking are enabled by entering the following command for each HBA. <pre>lsattr -El fscsin</pre> <i>n</i> is the controller/HBA ID. Output Example: The following output indicates that Dynamic Tracking is on (<i>dyntrk</i> value=yes) and the FastFail feature is enabled (<i>fc_err_recov</i>=fast_fail). <pre>attach switch How this adapter is CONNECTED False dyntrk yes Dynamic Tracking of FC Devices True fc_err_recov fast_fail FC Fabric Event Error RECOVERY Policy True scsi_id 0x10400 Adapter SCSI ID False sw_fc_class 3 FC Class for Fabric True</pre>

Step	Action																																																																	
10	Verify that the ASL is installed by entering the following command. vxddladm listsupport Result: The libvxnetapp.so library is listed in the command output. See the example below: <table><tr><td>LIBNAME</td><td>VID</td><td>PID</td></tr><tr><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>libvxf700.so</td><td>IBM</td><td>1742, 3552, 3542</td></tr><tr><td>libvxhds.so</td><td>HITACHI</td><td>All</td></tr><tr><td>libvxhitachi.so</td><td>HITACHI</td><td>DF350, DF400, DF400F</td></tr><tr><td>libvxpurple.so</td><td>SUN</td><td>T300</td></tr><tr><td>libvxshark.so</td><td>IBM</td><td>2105</td></tr><tr><td>libvxvpath.so</td><td>IBM</td><td>2105</td></tr><tr><td>libvxxp256.so</td><td>HP</td><td>All</td></tr><tr><td>libvxnetapp.so</td><td>NETAPP</td><td>All</td></tr></table>	LIBNAME	VID	PID	=====	=====	=====	libvxf700.so	IBM	1742, 3552, 3542	libvxhds.so	HITACHI	All	libvxhitachi.so	HITACHI	DF350, DF400, DF400F	libvxpurple.so	SUN	T300	libvxshark.so	IBM	2105	libvxvpath.so	IBM	2105	libvxxp256.so	HP	All	libvxnetapp.so	NETAPP	All																																			
LIBNAME	VID	PID																																																																
=====	=====	=====																																																																
libvxf700.so	IBM	1742, 3552, 3542																																																																
libvxhds.so	HITACHI	All																																																																
libvxhitachi.so	HITACHI	DF350, DF400, DF400F																																																																
libvxpurple.so	SUN	T300																																																																
libvxshark.so	IBM	2105																																																																
libvxvpath.so	IBM	2105																																																																
libvxxp256.so	HP	All																																																																
libvxnetapp.so	NETAPP	All																																																																
11	Enter the vxddmpadm listapm all command to verify that the APM is installed. Result: The output displays the Array Type as A/A-NETAPP. If the storage controller disks have not yet been detected by the host, the state of the APM is Not-Active. See the output example below. <table><tr><td>Module Name</td><td>APM Name</td><td>APM Version</td><td>Array Types</td><td>State</td></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>dmpaa</td><td>dmpaa</td><td>1</td><td>A/A</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P-C</td><td>Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-VERITAS</td><td>Not-Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-T3PLUS</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG-C</td><td>Not-Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>Disk</td><td>Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>APdisk</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/A-NETAPP</td><td>Not-Active</td></tr><tr><td>dmpsvc</td><td>dmpsvc</td><td>1</td><td>A/A-IBMSVC</td><td>Not-Active</td></tr></table>	Module Name	APM Name	APM Version	Array Types	State	=====	=====	=====	=====	=====	dmpaa	dmpaa	1	A/A	Active	dmpap	dmpap	1	A/P	Active	dmpap	dmpap	1	A/P-C	Active	dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active	dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active	dmpapg	dmpapg	1	A/PG	Not-Active	dmpapg	dmpapg	1	A/PG-C	Not-Active	dmpjbod	dmpjbod	1	Disk	Active	dmpjbod	dmpjbod	1	APdisk	Active	dmpnetapp	dmpnetapp	1	A/A-NETAPP	Not-Active	dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active
Module Name	APM Name	APM Version	Array Types	State																																																														
=====	=====	=====	=====	=====																																																														
dmpaa	dmpaa	1	A/A	Active																																																														
dmpap	dmpap	1	A/P	Active																																																														
dmpap	dmpap	1	A/P-C	Active																																																														
dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active																																																														
dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active																																																														
dmpapg	dmpapg	1	A/PG	Not-Active																																																														
dmpapg	dmpapg	1	A/PG-C	Not-Active																																																														
dmpjbod	dmpjbod	1	Disk	Active																																																														
dmpjbod	dmpjbod	1	APdisk	Active																																																														
dmpnetapp	dmpnetapp	1	A/A-NETAPP	Not-Active																																																														
dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active																																																														

Next steps

After you install the ASL and APM, complete the following procedures:

- ◆ If you have Data ONTAP 7.1 or higher and you upgraded to Attach Kit 1.3 from a previous version, it is recommended to change the cfmode setting of your clustered systems to single image and then reconfigure your host to discover the new paths to the disk. See “[Migrating a Storage Cluster to Single Image Mode](#)” on page 41.
- ◆ On the storage system, create LUNs and map them to igroups containing the WWPNs of the host HBAs. See “[Creating and mapping igroups and LUNs](#)” on page 57.
- ◆ On the host, discover the new LUNs and configure them to be managed by VxVM. See “[VERITAS software \(multipathing\)](#)” on page 62.

About this chapter

If you have a storage cluster, it is recommended that you migrate the `cfmode` setting of both systems in the cluster to single image mode. This chapter describes how to decide if you should migrate, as well as how to change the `cfmode` setting and reconfigure the host to discover new paths to its disks.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Deciding whether to migrate to single image mode”](#) on page 42
- ◆ [“Understanding `cfmode`”](#) on page 43
- ◆ [“Tasks to complete before you begin”](#) on page 47
- ◆ [“Upgrading to single image mode”](#) on page 49

Deciding whether to migrate to single image mode

When to migrate to single image mode

If you upgraded to Data ONTAP 7.1 or higher, it is recommended to change your storage systems' cfmodes setting to single image.

When you upgrade to Data ONTAP 7.1 from a previous release, your cfmodes settings are preserved. Your cfmodes setting is one of the following:

- ◆ FAS2xx models: dual-fabric mode
- ◆ All other storage systems and backup storage: partner or mixed.

To display the current cfmodes settings, enter the following command on each system in the cluster:

```
fcpx show cfmodes
```

The single image setting provides many advantages:

- ◆ It supports all storage systems and switches.
- ◆ It increases the number of available paths in your configuration.

For details, see “[Reasons for using single image mode](#)” on page 43.

Migrating to single image mode also requires planning downtime in your configuration, possible modifications to zoning, and reconfiguring your host to detect new paths to LUNs.

- ◆ If you want to migrate your cluster to single image mode, use the procedures in this chapter to change the cfmodes settings on your cluster and reconfigure your host.
- ◆ If you want to keep the current cfmodes setting for your cluster, proceed to Chapter 5, “[Creating, Configuring, and Managing LUNs](#),” on page 55.

Understanding cfmode

What is single image mode?

In single image mode, both systems in the cluster share the same World Wide Node Name (WWNN).

Single image mode provides the following advantages:

The host can access all LUNs through any target port on the cluster.

- ◆ Single image mode is supported on all storage systems and by all supported FCP hosts.
- ◆ There are no switch limitations. You can connect a storage system in single image mode to any supported FCP switch.

Reasons for using single image mode

You might want to migrate your configuration to single image mode for the following reasons:

- ◆ To increase the number of paths to the cluster
You can increase the number of available paths in your configuration because all ports on each storage system are available to the host.
- ◆ To support heterogeneous configurations
For example, you upgraded an existing SAN configuration with Solaris hosts and you want to add HP-UX or AIX hosts to the configuration.

HP-UX and AIX hosts do not support failover when the systems in a cluster are in standby mode. If you want to add these hosts to an existing homogenous configuration that has Solaris-only hosts, you have to change the cfmode setting on the cluster to either single image or partner.

Resolving LUN mapping conflicts

Single image mode does not allow the following mapping conflicts:

- ◆ Two LUNs, each one on a different node in the cluster, that have the same LUN ID and are mapped to the same igroup.
- ◆ Initiators that are members of more igroups with a different OS Type (for example, an initiator that is a member of a Solaris igroup on the local system and a member of an AIX igroup on the partner)

When you change the cfmode and restart the FCP service, one WWNN is automatically assigned to both systems in the cluster. This means you might have LUN mapping conflicts between the systems in the cluster.

Before you switch to the single image cfmode setting, you use the `lun config_check -S` command to check for any LUN mapping and ostype conflicts between the nodes in the cluster.

Impact of changing the cfmode setting

Changing the cfmode setting of a cluster impacts the following components of your configuration:

- ◆ Access to LUNs

When you change the cfmode setting, you change the available paths between the host and the storage cluster. Some previously available paths are no longer available and some new paths become available. You must reconfigure every host that is connected to the cluster to discover the new paths. The LUNs might be accessible but must not be used until you reconfigure the host.

- ◆ Multipathing policy

If you have multipathing software in your configuration, changing the cfmode setting might also affect the multipathing policy because there are more paths between the host and the cluster.

- ◆ Zoning


Updating the system to single image cfmode changes the WWNN of the cluster and the World Wide Port Names (WWPNs) of the target ports on each storage system. If the cluster connects to a switch by using soft zoning (zoning by WWPN), you must check the zoning setup and make the required updates to your switch. Systems that connect to switches using hard (physical port) zoning are not impacted.

- ◆ Switch configuration and cabling

In mixed and dual fabric modes, the target ports log into the fabric in loop mode. In single image mode, the target ports log into the fabric using point-to-point (PTP) mode. This means you might need to update your switch configuration when you migrate to single image mode so that the target connects to the fabric as point-to-point device.


Any currently supported cabling configuration is supported by single image mode. This means that if you have a supported configuration, it will continue to work after you migrate to single image mode. You might want to change the cabling of the cluster to the fabric to optimize high availability and performance. The following figures show examples of how to cable a cluster in single image mode to optimize high availability and performance.

Example cabling for single image mode
for systems with target HBAs



The FAS3000 storage systems have onboard Fibre Channel ports, which can be configured to operate in either target or initiator mode. In the following figure, ports 0b and 0d connect to the fabric and are configured in target mode. Ports 0a and 0c connect to the disk shelves and are configured in initiator mode. For information about configuring Fibre Channel ports on systems with onboard ports, see the *Data ONTAP Block Access Management Guide for FCP and iSCSI*.

Example cabling for single image mode
for systems with onboard FCP ports



For detailed information about the cabling configurations supported and recommended examples for systems with onboard FCP ports, see the *FCP Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Tasks to complete before you begin

Prerequisite tasks Complete the following tasks before you change your cfmode:

Task	Description	What to do
Save configuration information.	Enter commands that display information about the current configuration and save the output to the local disk. On the host, make a backup of all the volume groups that contain LUNs. These backups must contain up-to-date information for both the data within the volume group and the volume group configuration.	Perform a complete backup of the system's data. Save output from the following commands: ◆ <code>sanlun lun show</code> displays LUNs visible to the host. ◆ <code>sanlun lun show -p</code> displays LUNs and paths visible to the host If your LUNs are under VxVM disk management, then save the output from the following commands. ◆ <code>vxdisk list</code> displays disks under VxVM control ◆ <code>vxdisk path</code> displays disks and paths under VxVM control ◆ <code>vxprint htg diskgroup_name</code> displays VxVM volumes and status Note _____ If you are using the native disk manager, then use the appropriate commands to obtain information about your disks. _____
Identify affected hosts and initiators.	You must reconfigure and reboot any hosts and initiators that access LUNs on the systems you are upgrading.	Enter the following command: <pre>igroup show</pre>

Schedule downtime for the storage cluster.	Hosts cannot access LUNs while you change cfmode. You must allow sufficient downtime to change the cfmode setting and reconfigure and reboot the host.	Enter commands that quiesce host I/O and take host applications off-line.
--	--	---

Upgrading to single image mode

Procedure Overview

To migrate the host and storage cluster to single image mode you must complete the following tasks:

1. Prepare the host
Quiesce I/O and protect the configuration by unmounting file systems and disconnecting LUNs attached to the host. For example if you are using VxVM software, deport the VxVM disk groups.
2. Upgrade the storage cluster to single image mode
Resolve LUN mapping or OS-type conflicts, and change the cfmodesetting.
3. Validate zoning and cabling
Check zoning and resolve zone mapping issues on configurations that use soft zoning.
4. Reconfigure the host
Reconfigure the host to detect the new devices and new paths to the devices.

Preparing the host

To prepare the AIX host for to migrate to single image mode, complete the following steps.

Step	Action
1	Open a session to the AIX host.
2	Find all the LUNs and volume groups that contain LUNs from the cluster by entering the following command: sanlun lun show -p
3	Record the LUN volume group membership.
4	Unmount the file systems that contain the volume groups mapped to storage system LUNs. umount /filesystem_path
5	Quiesce the volume groups by entering the following command: varyoffvg volume_group_name
6	Export the volume group by entering the following command: exportvg volume_group_name
7	Enter the following command to list local devices: lsdev -Cc disk
8	Remove the existing NetApp LUN handles for the devices by entering the following command for each device: rmdev -dl device_name Example: rmdev -dl hdisk2

Step	Action
9	Verify that the device handles are removed by entering the following command: <pre>lsdev -Cc disk</pre> Example: The following example output shows that the devices handles were successfully removed because devices that were mapped to the LUNs are not listed. In this case, only internal devices are listed. <pre>hdisk0 Available 1S-08-00-8,0 16 Bit LVD SCSI Disk Drive hdisk1 Available 1S-08-00-9,0 16 Bit LVD SCSI Disk Drive</pre>
10	Proceed to “ Migrating the cluster to single image mode ” below.

Migrating the cluster to single image mode

To migrate the cluster to single image mode, complete the following steps.

Step	Action
1	Open a window to the console of each storage system you want to update.
2	Login at the console of the system(s) you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts. <pre>lun config_check -S</pre> Result: Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique to each node in the cluster. You use the <code>igroup</code> command to change igroup otypes, if required.
3	Switch to advanced privileges by entering the following command: <pre>priv set advanced</pre>
4	Stop the FCP service on both systems. On each system, enter: <pre>fcv stop</pre>

Step	Action
5	Set the cfmode to single_image. On each system console, enter: <pre>fcv set cfmode single_image</pre> If you see messages indicating mapping conflicts, then return to Step 2 to resolve the mapping conflicts.
6	Start the FCP service on each system in the cluster. On each system, enter the following command: <pre>fcv start</pre>
7	After you switch to the single image setting, both nodes in the cluster use the same WWNN. Obtain the WWNN of the cluster by entering the following command on the console of each system: <pre>fcv nodename</pre>
8	If you are using soft zoning, you need the WWPNNs of the target ports to reconfigure zones on your switch. To obtain the WWPNNs, enter the following command on each system in the cluster. <pre>fcv config</pre> For more information about zoning, see “Validating zoning and connectivity” on page 52.
9	Return to administrative privileges. On each storage system, enter: <pre>priv set admin</pre>

Validating zoning and connectivity

Changes to WWNNs and WWPNNs: After you change the cfmode setting to single image both systems in the cluster share a single WWNN. One node in the cluster takes the WWNN of its partner. This means that the WWNN of one of the systems in the cluster has changed. The WWPNNs of the target ports also change. You might have to update your zoning configuration depending on the type of zoning you use. Soft zoning uses the WWPNN. Port zoning uses the physical location (switch port) in the zone. If you use soft zoning you must update your zones to use the new WWPNNs. Configurations that use hard zoning are not impacted.

Cabling and zoning: If you did not change your cabling, migrating to single image mode also changes the number of available paths to the LUNs. Some additional paths to the target become available because more target ports are

visible to the host. You can use zoning or Data ONTAP portsets to reduce the number of paths in your configuration to comply with the number of paths supported by your multipathing software.

For details about validating and modifying switch zones see the documentation provided by your switch manufacturer.

For details about portsets, see the [Data ONTAP 7.1 Block Access Management Guide](#).

You might also want to change your cabling if you have more paths between the host and the cluster and want to reduce the number of paths. It is recommended that after you change the `cfmode` setting, make sure each HBA has one primary and at least one secondary path to each LUN. For additional information, see the FCP Configuration Guide at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Reconfiguring AIX hosts

To reconfigure AIX hosts to detect new paths to the storage, complete the following steps.

Step	Action
1	Connect to the AIX host console.
2	On the host console, enter the following command: cfgmgr The host scans the bus for new devices. If you created volume groups with the <code>autovaryon</code> flag, then the volume groups are reimported.
3	Verify the new paths by entering the following command. sanlun lun show -p
4	If you did not create the volume groups with the <code>autovaryon</code> flag, enter the following command: importvg -y volume_group_name disk_handle <i>volume_group_name</i> is the name you assigned to the volume group. <i>disk_handle</i> is the local device MPIO uses as the reference point for all paths to a LUN.
5	Enter the following command to verify the paths to the LUNs. sanlun lun show -p

Step	Action
6	Verify that the volume groups are active by entering the following command. lsvg -o
7	If you created file systems on the devices, mount the file systems.

About this chapter

This chapter describes how prepare the storage system for LUN configuration, and also how to create and configure LUNs.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Preparing the storage system”](#) on page 56
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 57
- ◆ [“LUN configuration overview”](#) on page 62
- ◆ [“Configuring LUNs with native AIX LVM”](#) on page 64
- ◆ [“Configuring LUNs with VxVM and DMP multipathing software”](#) on page 77
- ◆ [“Displaying LUN and HBA information”](#) on page 83
- ◆ [“Using sanlun to display host LUN information”](#) on page 84
- ◆ [“Displaying host HBA information”](#) on page 87

Preparing the storage system

Correct media type for direct-attached storage system

If you use a direct-attached configuration you must set the target HBA media type to loop. Use the `fcp config` command to set the media type.

Note_____

If you are using a cluster, the media type setting is disabled.

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

Setting the media type: To set the media type to loop, perform the following steps:

Step	Action
1	Stop each adapter used in the direct attach configuration by entering the following command: <code>fcp config <adapter> down</code> <i>adapter</i> is the name of the Fibre Channel target adapter.
2	Set the media type to loop by entering the following command for each adapter: <code>fcp config <adapter> mediatype loop</code>
3	Restart the adapter by entering the following command: <code>fcp config <adapter> up</code>
4	Reboot the storage system and the host.
5	Verify the adapter ports are online by entering the following command: <code>fcp show adapter</code>
6	Verify that the media type is now correctly set on each adapter port by entering the following command; <code>fcp config</code>

Creating and mapping igroups and LUNs

Tasks needed to create and map LUNs

Before you can complete the configuration, you must perform the following tasks:

- ◆ Create an igroup. An igroup is a collection of host WWPNs.
- ◆ Create one or more LUNs and map the LUNs to an igroup.

Best practices for creating LUNs

The best practices for creating LUNs include the following:

- ◆ Each igroup should map to an application and should include all the initiators used by the application to access its data. (Multiple applications may be using the same initiators.)
- ◆ Do not put LUNs in /vol/vol0 (the root volume).
- ◆ Ensure that the LUNs are at the root of a qtree in a volume.
- ◆ Enable space reservations.
- ◆ If you have a clustered pair, use different LUN IDs on each system in the cluster.
- ◆ Disable scheduled auto snapshots.

For more information about creating and managing LUNs, see the *Data ONTAP SAN System Administrator's Guide*.

Creating an igroup

To create an igroup, complete the following steps. If you have a cluster, you must create a separate igroup on each system in the cluster.

Note

The examples used in these steps assume that the application is running on a single host, not a cluster. They also assume that you have already attached the host and cluster to fabric.

Step	Action
1	Log in as root on the host.

Step	Action
2	Get the WWPN for each of the host's HBAs by entering the following command: <pre>sanlun fcp show adapter -c</pre> The <code>-c</code> option tells the <code>sanlun</code> utility to generate the output that is used to create the igroup. You take the resulting command line and run it on the storage system to create an igroup containing all the HBAs on your host. You can cut and paste the command from the host window to a storage system console window. Example: Running <code>sanlun fcp show adapter -c</code> produces the following output on an AIX system called <code>hostA</code>. This example uses two WWPNs. <pre># sanlun fcp show adapter -c</pre> Enter this filer command to create an initiator group for this system: <pre>igroup create -f -t aix "hostA" 10000000c930e3bd 10000000c930e2d8 #</pre> Because the name of the host is "hostA," this command line produces an igroup called "hostA."
3	On the host in a different session, use the <code>telnet</code> command to access the storage system.
4	Copy the <code>igroup create</code> command from Step 2. Example: If you use the information in Step 2, you would copy the following command line: <pre>igroup create -f -t aix "hostA" 10000000c930e3bd 10000000c930e2d8</pre>
5	Paste the command on the storage system's command line and press Enter. Result: The igroup is created with the same name as the host name.

Step	Action
6	On the storage system command line, verify that the new igroup has been correctly set up with AIX as the operating system by entering the command: <pre>igroup show -v</pre> Result: The storage system displays the newly created igroup with the host's WWPNs. Example: The following is an sample of the output from the <code>igroup show</code> command. Note If you want to change the igroup ostype to aix after you've created the igroup, you must manually set this attribute. See the chapter “Troubleshooting” on page 89 for information about how to do this. <pre>targetX> igroup show -v hostA hostA (FCP) : OS Type: aix Member: 10:00:00:00:c9:30:e3:bd Member: 10:00:00:00:c9:30:e2:d8 targetX></pre>

Creating and mapping LUNs

After you create a LUN, map it to the igroup you created in [“Creating an igroup”](#) on page 57. Complete the following steps.

To create a LUN and map it to an igroup, complete the following steps:

Step	Action
1	Enter the following command to create a LUN. <pre>lun create -s <i>size</i> -t <i>ostype</i> <i>lun-path</i></pre> <i>size</i> indicates the size of the LUN to be created, in bytes by default. You can use one of the following characters to multiply the size: c (bytes), w (words or double bytes), b (512), k (kilobytes), m (megabytes), g (gigabytes), t (terabytes). -t <i>ostype</i> indicates the type of operating system running on the host that accesses the boot LUN. Enter <i>aix</i>. <i>lun-path</i> is the LUN's path name that includes the volume and qtree.
2	Enter the following command to map the LUN to the igroup. <pre>lun map <i>lun-path</i> <i>igroup-name</i> <i>LUN-ID</i></pre> <i>lun-path</i> is the path name of the LUN you created. <i>igroup-name</i> is the name of the initiator group you created in “Creating an igroup” on page 57. <i>LUN-ID</i> is the identification number that you specify. The host associates the LUN ID with the location and path name of the LUN. If you do not specify a LUN ID, Data ONTAP automatically assigns one.
3	Verify that the LUN is mapped to the igroup by entering the following command: <pre>lun show -v <i>lun-path</i></pre> <i>lun-path</i> is the path to a specific LUN. Example: The following example verifies the LUN at vol/aix/10. <pre>targetX> lun show -v vol/aix/10 /vol/aix/10 10g (10737418240) (r/w, online, mapped)</pre>

For detailed information about creating and managing LUNs, see the *Block Access Management Guide* for your version of Data ONTAP.

Verifying the mapped LUN

To verify the mapped LUN, complete the following step.

Step	Action
1	On the command line, use the <code>lun show</code> command to display information about the mapped LUN. <pre>lun show -v lun-path</pre> <i>lun-path</i> is the path to a specific LUN. Result: The information for the specified LUN is displayed. Example: The following example verifies the LUN at vol/aix/10. <pre>targetX> lun show -v vol/aix/10 /vol/aix/10 10g (10737418240) (r/w, online, mapped) targetX></pre>

Next Steps

The next configuration task you must complete involves configuring the LUNs. LUN configuration procedures vary, and depend on your multipathing requirements and the software you choose.

- ◆ If your configuration has no multipathing requirements, refer to [“Configuring LUNs with native AIX LVM”](#) on page 64
- ◆ If your configuration includes multipathing with VERITAS software, refer to [“Configuring LUNs with VxVM and DMP multipathing software”](#) on page 77
- ◆ If you are unsure of your configuration and software, refer to [“LUN configuration overview”](#) on page 62

LUN configuration overview

Types of configurations and software

The software you use to configure LUNs on the host depends on your multipathing requirements. Systems that do not require multipathing can use Native AIX Logical Volume Manager (LVM).

Systems that require multipathing can use VxVM and DMP.

Note

Before you configure LUNs to be managed by VxVM, make sure you installed the ASL 4.0 and APM according to the procedures in Chapter 3, “[Installing the ASL and APM](#),” on page 35.

Configuring LUNs with native AIX LVM (no multipathing)

To configure LUNs using native AIX LVM:

1. Use `cfgmgr` to get the host to discover the new LUNs.
2. Create a volume group.
3. Configure data access (raw, jfs, or jfs2 file systems).

Refer to “[Configuring LUNs with native AIX LVM](#)” on page 64 for additional information.

VERITAS software (multipathing)

To configure LUNs using VxVM and DMP:

1. Set the `vxrestore` daemon interval to 60 seconds.
2. Use `cfgmgr` to get the host to discover the new LUNs.
3. Use `vxctl enable` to allow VxVM to recognize AIX hdisks.
4. Use VxVM tools (VEA, `vxdiskadm`, CLI) or `smit` to add disks to the VxVM disk group and initialize them for use by VxVM.
5. Use VxVM tools (VEA, CLI) or `smit` to create a volume

6. Use VxVM tools (VEA,CLI) or `smit` to configure data access (raw or file system)

Refer to “[Configuring LUNs with VxVM and DMP multipathing software](#)” on page 77 for additional information.

Configuring the LUNs

Configuring LUNs with native AIX LVM

Getting the host to discover the new LUNs

To get the host to discover the new LUNs, complete the following steps.

Note_____

You do not need to reboot the host to discover the LUNs.

Step	Action
1	Log in as root on the host.
2	Execute the <code>cfgmgr</code> command on the AIX host to discover the newly created LUNs: # <code>cfgmgr</code> Result: The <code>cfgmgr</code> command causes the AIX system to: ◆ Log into the fabric and to check for new devices. ◆ Create new device entries.

Step	Action
3	Enter <code>sanlun lun show</code> to verify that the host has discovered the LUNs. Example: This example shows typical output of <code>sanlun lun show</code> that for a direct-attached configuration. <pre># sanlun lun show filer: lun-pathname device filename adapter lun size lun state systemX: /vol/aix/hostA-systemX-lun0 hdisk2 fcs0 200m (209715200) GOOD systemX: /vol/aix/hostA-systemX-lun1 hdisk3 fcs0 1.2g (1258291200) GOOD systemX: /vol/aix/hostA-systemX-lun2 hdisk4 fcs0 2.1g (2306867200) GOOD systemX: /vol/aix/hostA-systemX-lun3 hdisk5 fcs0 3.1g (3355443200) GOOD systemY: /vol/aix/hostA-systemY-lun4 hdisk6 fcs1 500m (524288000) GOOD systemY: /vol/aix/hostA-systemY-lun5 hdisk7 fcs1 2g (2147483648) GOOD systemY: /vol/aix/hostA-systemY-lun6 hdisk8 fcs1 1g (1073741824) GOOD systemY: /vol/aix/hostA-systemY-lun7 hdisk9 fcs1 3g (3221225472) GOOD systemY: /vol/aix/hostA-systemY-lun8 hdisk10 fcs1 4g (4294967296) GOOD</pre>

Creating a volume group

You can use whichever tool you normally use to create volume groups. The following steps use SMIT.

Note

See the IBM AIX documentation for information on managing volume groups, logical volumes and file systems.

Step	Action
1	Use the AIX System Management Interface Tool (SMIT) to create a volume group by entering the following command: <pre>smit vg</pre>

Step	Action																																										
2	Specify a volume group name, and enter the information to set up the volume group at the SMIT Add a Volume Group screen. Add a Volume Group Type or select values in entry fields. Press Enter AFTER making all desired changes. <table><tr><td></td><td>[Entry Fields]</td><td></td></tr><tr><td>VOLUME GROUP name</td><td>[newvg]</td><td></td></tr><tr><td>Physical partition SIZE in megabytes</td><td>4</td><td>+</td></tr><tr><td>* PHYSICAL VOLUME names</td><td>[hdisk6]</td><td>+</td></tr><tr><td>Activate volume group AUTOMATICALLY at system restart?</td><td>yes</td><td>+</td></tr><tr><td>Volume Group MAJOR NUMBER</td><td>[]</td><td>+#</td></tr><tr><td>Create VG Concurrent Capable?</td><td>no</td><td>+</td></tr><tr><td>Auto-varyon in Concurrent Mode?</td><td>no</td><td>+</td></tr><tr><td>LTG Size in kbytes</td><td>128</td><td>+</td></tr><tr><td>LTG Size in kbytes</td><td>128</td><td>=Help</td></tr><tr><td>F2=Refresh</td><td>F3=Cancel</td><td>F4=List</td></tr><tr><td>F5=Reset</td><td>F6=Command</td><td>F7=Edit</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td><td>Enter=Do</td></tr><tr><td>F8=Image</td><td></td><td></td></tr></table> -----		[Entry Fields]		VOLUME GROUP name	[newvg]		Physical partition SIZE in megabytes	4	+	* PHYSICAL VOLUME names	[hdisk6]	+	Activate volume group AUTOMATICALLY at system restart?	yes	+	Volume Group MAJOR NUMBER	[]	+#	Create VG Concurrent Capable?	no	+	Auto-varyon in Concurrent Mode?	no	+	LTG Size in kbytes	128	+	LTG Size in kbytes	128	=Help	F2=Refresh	F3=Cancel	F4=List	F5=Reset	F6=Command	F7=Edit	F9=Shell	F10=Exit	Enter=Do	F8=Image		
	[Entry Fields]																																										
VOLUME GROUP name	[newvg]																																										
Physical partition SIZE in megabytes	4	+																																									
* PHYSICAL VOLUME names	[hdisk6]	+																																									
Activate volume group AUTOMATICALLY at system restart?	yes	+																																									
Volume Group MAJOR NUMBER	[]	+#																																									
Create VG Concurrent Capable?	no	+																																									
Auto-varyon in Concurrent Mode?	no	+																																									
LTG Size in kbytes	128	+																																									
LTG Size in kbytes	128	=Help																																									
F2=Refresh	F3=Cancel	F4=List																																									
F5=Reset	F6=Command	F7=Edit																																									
F9=Shell	F10=Exit	Enter=Do																																									
F8=Image																																											
3	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. 0516-1254 mkvg: Changing the PVID in the ODM. newvg <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr><tr><td>n=Find Next</td><td></td><td></td><td></td></tr><tr><td><F10 exits smit></td><td></td><td></td><td></td></tr></table> -----	F1=Help	F2=Refresh	F3=Cancel	F6=Command	F8=Image	F9=Shell	F10=Exit	/=Find	n=Find Next				<F10 exits smit>																													
F1=Help	F2=Refresh	F3=Cancel	F6=Command																																								
F8=Image	F9=Shell	F10=Exit	/=Find																																								
n=Find Next																																											
<F10 exits smit>																																											
4	Press F10 to exit SMIT.																																										

Step	Action
5	Check the information on the new volume group. <pre># lsvg newvg VOLUME GROUP: newvg VG IDENTIFIER: 000f3bca00004c00000000f8267ba452 VG STATE: active PP SIZE: 16 megabyte(s) VG PERMISSION: read/write TOTAL PPs: 639 (10224 megabytes) MAX LVs: 256 FREE PPs: 639 (10224 megabytes) LVs: 0 USED PPs: 0 (0 megabytes) OPEN LVs: 0 QUORUM: 2 TOTAL PVs: 1 VG DESCRIPTORS: 2 STALE PVs: 0 STALE PPs: 0 ACTIVE PVs: 1 AUTO ON: yes MAX PPs per PV: 1016 MAX PVs: 32 LTG size: 128 kilobyte(s) AUTO SYNC: no HOT SPARE: no</pre>

Access storage on a volume group

You can access storage space on a volume group using any of the methods you normally use. These can include using a raw logical volume, jfs, or jfs2.

Create a file system

You can use whichever tool you normally use to create file system, if you need one. The following steps use SMIT.

Note

See the IBM documentation for information on managing file systems.

Step	Action
1	To use SMIT to set up a file system, enter the following command: <pre>smit fs</pre>

Step	Action
2	Enter the information to set up the file system at the SMIT Add a Journaled File System screen. <pre> ----- -->Add / Change / Show / Delete File Systems-->Journaled File Systems--> Add a - --> Journaled File System -->Add a Standard Journaled File System Add a Journaled File System Move cursor to desired item and press Enter. Add a Standard Journaled File System Add a Compressed Journaled File System Add a Large File Enabled Journaled File System +-----+ Volume Group Name Move cursor to desired item and press Enter. rootvg newvg F1=Help F2=Refresh F3=Cancel F8=Image F10=Exit Enter=Do F1=H /=Find n=Find Next F9=S +-----+ </pre>

Step	Action																																													
3	Select the volume group you created in the section “Creating a volume group” on page 65. ----- Add a Standard Journaled File System Type or select values in entry fields. Press Enter AFTER making all desired changes. <table><tr><td></td><td>[Entry Fields]</td><td></td></tr><tr><td>Volume group name</td><td>newvg</td><td></td></tr><tr><td>* SIZE of file system (in 512-byte blocks)</td><td>[2140000]</td><td>#</td></tr><tr><td>* MOUNT POINT</td><td>[/newfs]</td><td></td></tr><tr><td>Mount AUTOMATICALLY at system restart?</td><td>yes</td><td>+</td></tr><tr><td>PERMISSIONS</td><td>read/write</td><td>+</td></tr><tr><td>Mount OPTIONS</td><td>[]</td><td>+</td></tr><tr><td>Start Disk Accounting?</td><td>no</td><td>+</td></tr><tr><td>Fragment Size (bytes)</td><td>4096</td><td>+</td></tr><tr><td>Number of bytes per inode</td><td>4096</td><td>+</td></tr><tr><td>Allocation Group Size (MBytes)</td><td>8</td><td>+</td></tr></table> <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F4=List</td></tr><tr><td>F5=Reset</td><td>F6=Command</td><td>F7=Edit</td><td>F8=Image</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td><td>Enter=Do</td><td></td></tr></table> -----		[Entry Fields]		Volume group name	newvg		* SIZE of file system (in 512-byte blocks)	[2140000]	#	* MOUNT POINT	[/newfs]		Mount AUTOMATICALLY at system restart?	yes	+	PERMISSIONS	read/write	+	Mount OPTIONS	[]	+	Start Disk Accounting?	no	+	Fragment Size (bytes)	4096	+	Number of bytes per inode	4096	+	Allocation Group Size (MBytes)	8	+	F1=Help	F2=Refresh	F3=Cancel	F4=List	F5=Reset	F6=Command	F7=Edit	F8=Image	F9=Shell	F10=Exit	Enter=Do	
	[Entry Fields]																																													
Volume group name	newvg																																													
* SIZE of file system (in 512-byte blocks)	[2140000]	#																																												
* MOUNT POINT	[/newfs]																																													
Mount AUTOMATICALLY at system restart?	yes	+																																												
PERMISSIONS	read/write	+																																												
Mount OPTIONS	[]	+																																												
Start Disk Accounting?	no	+																																												
Fragment Size (bytes)	4096	+																																												
Number of bytes per inode	4096	+																																												
Allocation Group Size (MBytes)	8	+																																												
F1=Help	F2=Refresh	F3=Cancel	F4=List																																											
F5=Reset	F6=Command	F7=Edit	F8=Image																																											
F9=Shell	F10=Exit	Enter=Do																																												
4	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. Based on the parameters chosen, the new /newfs JFS file system is limited to a maximum size of 134217728 (512 byte blocks) New File System size is 2162688 <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr><tr><td>n=Find Next</td><td></td><td></td><td></td></tr></table> -----	F1=Help	F2=Refresh	F3=Cancel	F6=Command	F8=Image	F9=Shell	F10=Exit	/=Find	n=Find Next																																				
F1=Help	F2=Refresh	F3=Cancel	F6=Command																																											
F8=Image	F9=Shell	F10=Exit	/=Find																																											
n=Find Next																																														

Step	Action														
5	Press F10 to exit SMIT.														
6	Mount the new file system. # mount /newfs # df -k /newfs <table><tr><td>Filesystem</td><td>1024-blocks</td><td>Free</td><td>%Used</td><td>Iused</td><td>%Iused</td><td>Mounted on</td></tr><tr><td>/dev/lv01</td><td>1081344</td><td>1047356</td><td>4%</td><td>17</td><td>1%</td><td>/newfs</td></tr></table>	Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on	/dev/lv01	1081344	1047356	4%	17	1%	/newfs
Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on									
/dev/lv01	1081344	1047356	4%	17	1%	/newfs									

Determine which device to use as the physical volume

You must determine which device to use as the physical volume for the Logical Volume Manager (LVM) before you create a volume group. You can use the SAN Toolkit's `sanlun lun show -p` command to determine which host disk device VERITAS DMP uses as the I/O access device for the LUN.

Using sanlun to find the access device: The following example shows part of the output of the `sanlun lun show -p` command. This example has been shortened.

```
# sanlun lun show -p
NetApp_PATH: shfas3020-07:/vol/aix/TestLUNd
LUN: 3
LUN Size:      1g (1073741824)
Host Device: hdisk5
LUN State: GOOD
Controller_CF_State: Cluster Enabled
Controller Partner: 1a
Multipath Provider: VERITAS
Multipathing Algorithm: round_robin
```

```

-----
VERITAS Controller AIX                      Controller VERITAS
path      path  VERITAS host      target HBA  path
status    type          path  HBA    port      priority
-----
Enabled    primary path0  fscsi0    1a        1
Enabled    primary path1  fscsi0    1b        1
Enabled    secondary path2  fscsi0    1a        1
Enabled    secondary path3  fscsi0    1b        1
Enabled    primary path4    fscsi1    1a        1
Enabled    primary path5    fscsi1    1b        1
Enabled    secondary path6   fscsi1    1a        1
Enabled    secondary path7   fscsi1    1b        1

```

```

NetApp_PATH: shfas3020-07:/vol/aix/TestLUNb
LUN: 1
LUN Size:    10g (10737418240)
Host Device: hdisk3
LUN State: GOOD
Controller_CF_State: Cluster Enabled
Controller Partner: 1a
Multipath Provider: AIX Native
Multipathing Algorithm: round_robin

```

```

-----
VERITAS Controller AIX                      Controller VERITAS
path      path  VERITAS host      target HBA  path
status    type          path  HBA    port      priority
-----
Enabled    primary path0  fscsi0    1a        1
Enabled    primary path1  fscsi0    1b        1
Enabled    secondary path2  fscsi0    1a        1
Enabled    secondary path3  fscsi0    1b        1
...

```

Creating a volume group

You can use whichever tool you normally use to create volume groups. The following steps use SMIT.

Step	Action
1	Determine which device to use as the physical volume for the LVM. See the previous section, “Determine which device to use as the physical volume” on page 70 for information on doing this.

Step	Action																																																								
2	Use the AIX System Management Interface Tool (SMIT) to create a volume group by entering the following command: <pre>smit vg</pre>																																																								
3	Enter a volume group name, and the information to set up the volume group at the SMIT Add a Volume Group screen. Add a Volume Group Type or select values in entry fields. Press Enter AFTER making all desired changes. <table><tr><td colspan="2"></td><td>[Entry Fields]</td><td></td></tr><tr><td>VOLUME GROUP name</td><td></td><td>[newvg]</td><td></td></tr><tr><td>Physical partition SIZE in megabytes</td><td></td><td>4</td><td>+</td></tr><tr><td>* PHYSICAL VOLUME names</td><td></td><td>[hdisk6]</td><td>+</td></tr><tr><td>Activate volume group AUTOMATICALLY at system restart?</td><td></td><td>yes</td><td>+</td></tr><tr><td>Volume Group MAJOR NUMBER</td><td></td><td>[]</td><td>+#</td></tr><tr><td>Create VG Concurrent Capable?</td><td></td><td>no</td><td>+</td></tr><tr><td>Auto-varyon in Concurrent Mode?</td><td></td><td>no</td><td>+</td></tr><tr><td>LTG Size in kbytes</td><td></td><td>128</td><td>+</td></tr><tr><td colspan="2">LTG Size in kbytes</td><td>128</td><td>=Help</td></tr><tr><td>F2=Refresh</td><td>F3=Cancel</td><td>F4=List</td><td></td></tr><tr><td>F5=Reset</td><td>F6=Command</td><td>F7=Edit</td><td>F8=Image</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td><td>Enter=Do</td><td></td></tr><tr><td colspan="4">-----</td></tr></table>			[Entry Fields]		VOLUME GROUP name		[newvg]		Physical partition SIZE in megabytes		4	+	* PHYSICAL VOLUME names		[hdisk6]	+	Activate volume group AUTOMATICALLY at system restart?		yes	+	Volume Group MAJOR NUMBER		[]	+#	Create VG Concurrent Capable?		no	+	Auto-varyon in Concurrent Mode?		no	+	LTG Size in kbytes		128	+	LTG Size in kbytes		128	=Help	F2=Refresh	F3=Cancel	F4=List		F5=Reset	F6=Command	F7=Edit	F8=Image	F9=Shell	F10=Exit	Enter=Do		-----			
		[Entry Fields]																																																							
VOLUME GROUP name		[newvg]																																																							
Physical partition SIZE in megabytes		4	+																																																						
* PHYSICAL VOLUME names		[hdisk6]	+																																																						
Activate volume group AUTOMATICALLY at system restart?		yes	+																																																						
Volume Group MAJOR NUMBER		[]	+#																																																						
Create VG Concurrent Capable?		no	+																																																						
Auto-varyon in Concurrent Mode?		no	+																																																						
LTG Size in kbytes		128	+																																																						
LTG Size in kbytes		128	=Help																																																						
F2=Refresh	F3=Cancel	F4=List																																																							
F5=Reset	F6=Command	F7=Edit	F8=Image																																																						
F9=Shell	F10=Exit	Enter=Do																																																							

4	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. 0516-1254 mkvg: Changing the PVID in the ODM. newvg <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr></table> n=Find Next <F10 exits smit> -----	F1=Help	F2=Refresh	F3=Cancel	F6=Command	F8=Image	F9=Shell	F10=Exit	/=Find																																																
F1=Help	F2=Refresh	F3=Cancel	F6=Command																																																						
F8=Image	F9=Shell	F10=Exit	/=Find																																																						

Step	Action
5	Press F10 to exit SMIT.
6	Check the information on the new volume group. <pre># lsvg newvg VOLUME GROUP: newvg VG IDENTIFIER: 000f3bca00004c00000000f8267ba452 VG STATE: active PP SIZE: 16 megabyte(s) VG PERMISSION: read/write TOTAL PPs: 639 (10224 megabytes) MAX LVs: 256 FREE PPs: 639 (10224 megabytes) LVs: 0 USED PPs: 0 (0 megabytes) OPEN LVs: 0 QUORUM: 2 TOTAL PVs: 1 VG DESCRIPTORS: 2 STALE PVs: 0 STALE PPs: 0 ACTIVE PVs: 1 AUTO ON: yes MAX PPs per PV: 1016 MAX PVs: 32 LTG size: 128 kilobyte(s) AUTO SYNC: no HOT SPARE: no</pre>

Access storage on a volume group

You can access storage space on a volume group using any of the methods you normally use. These can include using jfs, jfs2, or a raw logical volume.

Create a file system

You can use whichever tool you normally use to create a file system if you need one. The following steps use SMIT.

Note

See the IBM documentation for information on managing file systems.

Step	Action
1	To use SMIT to set up a file system, enter the following command: <pre>smit fs</pre>

Step	Action
2	Enter the information to set up the file system at the SMIT Add a Journaled File System screen. <pre> ----- -->Add / Change / Show / Delete File Systems-->Journaled File Systems--> Add a - --> Journaled File System -->Add a Standard Journaled File System Add a Journaled File System Move cursor to desired item and press Enter. Add a Standard Journaled File System Add a Compressed Journaled File System Add a Large File Enabled Journaled File System +-----+ Volume Group Name Move cursor to desired item and press Enter. rootvg newvg F1=Help F2=Refresh F3=Cancel F8=Image F10=Exit Enter=Do F1=H /=Find n=Find Next F9=S +-----+ </pre>

Step	Action																																													
3	Select the volume group you created in the section “Creating a volume group” on page 65. ----- Add a Standard Journaled File System Type or select values in entry fields. Press Enter AFTER making all desired changes. <table><tr><td></td><td>[Entry Fields]</td><td></td></tr><tr><td>Volume group name</td><td>newvg</td><td></td></tr><tr><td>* SIZE of file system (in 512-byte blocks)</td><td>[2140000]</td><td>#</td></tr><tr><td>* MOUNT POINT</td><td>[/newfs]</td><td></td></tr><tr><td>Mount AUTOMATICALLY at system restart?</td><td>yes</td><td>+</td></tr><tr><td>PERMISSIONS</td><td>read/write</td><td>+</td></tr><tr><td>Mount OPTIONS</td><td>[]</td><td>+</td></tr><tr><td>Start Disk Accounting?</td><td>no</td><td>+</td></tr><tr><td>Fragment Size (bytes)</td><td>4096</td><td>+</td></tr><tr><td>Number of bytes per inode</td><td>4096</td><td>+</td></tr><tr><td>Allocation Group Size (MBytes)</td><td>8</td><td>+</td></tr></table> <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F4=List</td></tr><tr><td>F5=Reset</td><td>F6=Command</td><td>F7=Edit</td><td>F8=Image</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td><td>Enter=Do</td><td></td></tr></table> -----		[Entry Fields]		Volume group name	newvg		* SIZE of file system (in 512-byte blocks)	[2140000]	#	* MOUNT POINT	[/newfs]		Mount AUTOMATICALLY at system restart?	yes	+	PERMISSIONS	read/write	+	Mount OPTIONS	[]	+	Start Disk Accounting?	no	+	Fragment Size (bytes)	4096	+	Number of bytes per inode	4096	+	Allocation Group Size (MBytes)	8	+	F1=Help	F2=Refresh	F3=Cancel	F4=List	F5=Reset	F6=Command	F7=Edit	F8=Image	F9=Shell	F10=Exit	Enter=Do	
	[Entry Fields]																																													
Volume group name	newvg																																													
* SIZE of file system (in 512-byte blocks)	[2140000]	#																																												
* MOUNT POINT	[/newfs]																																													
Mount AUTOMATICALLY at system restart?	yes	+																																												
PERMISSIONS	read/write	+																																												
Mount OPTIONS	[]	+																																												
Start Disk Accounting?	no	+																																												
Fragment Size (bytes)	4096	+																																												
Number of bytes per inode	4096	+																																												
Allocation Group Size (MBytes)	8	+																																												
F1=Help	F2=Refresh	F3=Cancel	F4=List																																											
F5=Reset	F6=Command	F7=Edit	F8=Image																																											
F9=Shell	F10=Exit	Enter=Do																																												
4	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. Based on the parameters chosen, the new /newfs JFS file system is limited to a maximum size of 134217728 (512 byte blocks) New File System size is 2162688 <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr><tr><td>n=Find Next</td><td></td><td></td><td></td></tr></table> -----	F1=Help	F2=Refresh	F3=Cancel	F6=Command	F8=Image	F9=Shell	F10=Exit	/=Find	n=Find Next																																				
F1=Help	F2=Refresh	F3=Cancel	F6=Command																																											
F8=Image	F9=Shell	F10=Exit	/=Find																																											
n=Find Next																																														

Step	Action														
5	Press F10 to exit SMIT.														
6	Mount the new file system. # mount /newfs # df -k /newfs <table><thead><tr><th>Filesystem</th><th>1024-blocks</th><th>Free</th><th>%Used</th><th>Iused</th><th>%Iused</th><th>Mounted on</th></tr></thead><tbody><tr><td>/dev/lv01</td><td>1081344</td><td>1047356</td><td>4%</td><td>17</td><td>1%</td><td>/newfs</td></tr></tbody></table>	Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on	/dev/lv01	1081344	1047356	4%	17	1%	/newfs
Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on									
/dev/lv01	1081344	1047356	4%	17	1%	/newfs									

Configuring LUNs with VxVM and DMP multipathing software

Setting the vxrestore daemon interval

The default policy for vxrestore daemon is to check disabled paths every 300 seconds. In the environment, you reduce this interval to 60 seconds to improve I/O failover handling. To set the vxrestore daemon, complete the following steps.

Step	Action
1	Stop the vxrestore daemon. <code>/usr/sbin/vxdmpadm stop restore</code>
2	Change the vxrestore daemon setting to 60. <code>/usr/sbin/vxdmpadm start restore policy=check_disabled interval=60</code> Note _____ This step reconfigures and restarts the vxrestore daemon without the need for an immediate reboot. _____
3	Edit the /etc/init.d/vxvm-sysboot file to make the new restore daemon interval persistent across reboots. By default, the restore daemon options are: <code>restore_daemon_opts="interval=300 policy=check_disabled"</code> Edit the restore daemon options so that the interval is 60 seconds: <code>restore_daemon_opts="interval=60 policy=check_disabled"</code>
4	Save and exit the /etc/init.d/vxvm-sysboot file.
5	Verify the changes. <code>/usr/sbin/vxdmpadm stat restored</code> Result: The command output shows the status of the vxrestore daemon, as shown by the output example below. The number of daemons running : 1 The interval of daemon: 60 The policy of daemon: check_disabled

Discovering the new LUNs

To get the host and VxVM to discover the new LUNs, complete the following steps.

Note

You do not need to reboot the host to discover the LUNs.

Step	Action																									
1	Log in as root on the host.																									
2	Execute the <code>cfgmgr</code> command on the AIX host to discover the newly created LUNs: <pre># cfgmgr</pre> Result: The <code>cfgmgr</code> command causes initiators to log into the fabric, the system to check for new devices, and the system to create new device entries.																									
3	Enter <code>lsdev -Cc disk</code> to verify that the host has discovered the LUNs. Example: This example shows the hdisks discovered by the AIX host. <pre>lsdev -Cc disk hdisk0 Available 08-08-00-5,0 16 Bit LVD SCSI Disk Drive hdisk1 Available 08-08-00-8,0 16 Bit LVD SCSI Disk Drive hdisk2 Available 04-08-02 VERITAS NetApp FCP Default PCM Disk hdisk3 Available 04-08-02 VERITAS NetApp FCP Default PCM Disk hdisk4 Available 04-08-02 VERITAS NetApp FCP Default PCM Disk hdisk5 Available 04-08-02 VERITAS NetApp FCP Default PCM Disk</pre>																									
4	Enter <code>vxctl enable</code> to enable VxVM to recognize each known hdisk. Result: The <code>vxctl</code> utility recognizes and assigns names to LUNs on the storage system. The following shows the results of <code>vxctl enable</code>. Example: The following <pre>vxdisk list</pre> <table><tr><th>DEVICE</th><th>TYPE</th><th>DISK</th><th>GROUP</th><th>STATUS</th></tr><tr><td>Disk_0</td><td>auto:LVM</td><td>-</td><td>-</td><td>LVM</td></tr><tr><td>Disk_1</td><td>auto:LVM</td><td>-</td><td>-</td><td>LVM</td></tr><tr><td>Netapp1_0</td><td>auto:none</td><td>-</td><td>-</td><td>online:invalid</td></tr><tr><td>Netapp1_0</td><td>auto:none</td><td>-</td><td>-</td><td>online:invalid</td></tr></table>	DEVICE	TYPE	DISK	GROUP	STATUS	Disk_0	auto:LVM	-	-	LVM	Disk_1	auto:LVM	-	-	LVM	Netapp1_0	auto:none	-	-	online:invalid	Netapp1_0	auto:none	-	-	online:invalid
DEVICE	TYPE	DISK	GROUP	STATUS																						
Disk_0	auto:LVM	-	-	LVM																						
Disk_1	auto:LVM	-	-	LVM																						
Netapp1_0	auto:none	-	-	online:invalid																						
Netapp1_0	auto:none	-	-	online:invalid																						

Step	Action
5	Verify that the ASL is has detected the new devices by entering the <code>vxddmpadm listenclosure all</code> command. In the output, verify that the model names of your storage systems are listed as the <code>ENCLR_NAME</code> and <code>ENCLR_TYPE</code>. Verify that the <code>ARRAY_TYPE</code> is <code>A/A NETAPP</code>. <pre> vxddmpadm listenclosure all ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE ===== Disk Disk DISKS CONNECTED Disk F9600 FAS960 1024916 CONNECTED A/A-NETAPP F9601 FAS960 1024915 CONNECTED A/A-NETAPP </pre>
6	Enter the <code>vxddmpadm listapm all</code> command to verify that the APM is active. Output example: The following output example shows that the NetApp APM is active. <pre> Module Name APM Name APM Version Array Types State ===== dmpaa dmpaa 1 A/A Active dmpap dmpap 1 A/P Active dmpap dmpap 1 A/P-C Active dmpapf dmpapf 1 A/PF-VERITAS Not-Active dmpapf dmpapf 1 A/PF-T3PLUS Not-Active dmpapg dmpapg 1 A/PG Not-Active dmpapg dmpapg 1 A/PG-C Not-Active dmpjbod dmpjbod 1 Disk Active dmpjbod dmpjbod 1 APdisk Active dmpnetapp dmpnetapp 1 A/A-NETAPP Active dmpsvc dmpsvc 1 A/A-IBMSVC Not-Active </pre>

Adding the disk to a disk group

Adding a disk to a disk group makes the disk space available when you create a VxVM volume. Disks that you add must be uninitialized, or free. If the disk is uninitialized, you must initialize it before you add it to the group.

Adding a disk to a disk group involves specifying how you want the to configure the disks. You should understand the following before you begin:

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

To add the disk to the group, complete the following step. Refer to VERITAS documentation for additional information.

Step	Action
1	Select any of the following methods: VEA: 1. Select the disk to add or the disk group to expand. 2. Select Actions and Add Disk to Dynamic Disk Group. 3. Indicate the disk group that you want. 4. Indicate the disks you want to add. vxdiskadm: 1. Enter vxdiskadm to activate the vxdiskadm interface. 2. Select option 1 Add or initialize one or more disks and follow screen directions. CLI: ◆ To configure private and public regions on a VxVM disk, enter vxdisksestap. ◆ To add one or more disks to a disk group, enter vx dg addisk smit: 1. Enter smit vxvmdiskadmadd to add disks using smit.

Creating a volume

VERITAS volumes span data across multiple disks. Creating a volume with VxVM involves selecting the disk group and disks, and also specifying the volume layout.

To create a VERITAS volume, complete the following step. Refer to VERITAS documentation for additional information.

Step	Action
1	Select any of the following methods: VEA: 1. Expand the disk group and select the disk group you want. 2. Select Actions and New Volume. 3. Specify the volume characteristics that meet your requirements. CLI: ◆ Enter the following: <code>vxassist -g <i>diskgroup</i> make <i>volume_name</i> <i>length</i> <i>attributes</i></code> ◆ To add one or more disks to a disk group, enter <code>vxvg addisk</code> smit: 1. Enter <code>smit vxvm.</code> 2. Select VxVM Volumes. 3. Select Add a VxVM Volume. 4. Select the layout that meets your requirements. 5. Select the volume characteristics.

Access storage on the VERITAS volume

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

To create a file system, follow the step below.

Step	Action
1	Select any of the following methods: VEA: 1. Select the volume that will contain the file system. 2. Specify Actions, File System, and New File System. 3. Specify the file system requirements. CLI: To create the file system, enter <pre>mount -V vxfs /dev/vx/dsk/<i>diskgroup/volume</i></pre> To mount the file system, enter <pre>mount -V vxfs /dev/vx/dsk/<i>diskgroup/volume</i> <i>mount_point</i></pre> smit: Enter <code>smit crvxfs1vstd</code>

Displaying LUN and HBA information

Utility for displaying information

The `sanlun` utility is a SAN Toolkit tool that aids in collecting and reporting information about paths to your devices and how they map to LUNs on the storage system. You can also use this command to display information about the host HBAs.

For detailed information

For detailed information, see the following topics:

- ◆ [“Using sanlun to display host LUN information”](#) on page 84
- ◆ [“Displaying host HBA information”](#) on page 87

Using sanlun to display host LUN information

Displaying host LUN information

To display host LUN information, complete the following steps.

Step	Action
1	Log in as root on the host.
2	At the host command line, use the <code>sanlun lun show</code> command to display LUN information: <pre>sanlun lun show [-v] [-d <i>host device filename</i> <i>all</i> <i>filer name</i> <i>filer name:filer pathname</i>]</pre> <code>-v</code> produces verbose output. <code>-d</code> is the device option and can be one of the following: ◆ <i>host device filename</i> specifies the special device file on the host. ◆ <i>all</i> lists all LUNs attached to the host. ◆ <i>filer name</i> is the host name of the storage system. ◆ <i>filer pathname</i> is the path name of the LUN on the storage system. (Multipathing only) If you want to see information about the primary and secondary paths, use the version of <code>sanlun lun show</code> with the <code>-p</code> option (you cannot use the <code>-d</code> option if you use <code>-p</code>): <pre>sanlun lun show -p [<i>filer name:filer pathname</i> <i>filer name</i> <i>all</i>]</pre> <code>-p</code> provides information about the primary and secondary paths available to the LUN when you are using multipathing. If you enter <code>sanlun lun show</code>, <code>sanlun lun show -p</code>, or <code>sanlun lun show -v</code> without any parameters, the utility responds as if you had included the <code>all</code> parameter.

Step	Action
	So you might enter: ◆ <code>sanlun lun show -p</code> 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. ◆ <code>sanlun lun show -d /dev/hdisk<x></code> to display the summary listing of the LUN(s) associated with the host device <code>/dev/hdisk<x></code> (<i>x</i> is the number of the hdisk; for example, you might have <code>hdisk1</code> or <code>hdisk4</code>). ◆ <code>sanlun lun show -v all</code> to display verbose output for all the LUN(s) currently available on the host. ◆ <code>sanlun lun show toaster</code> to display a summary listing of all the LUNs available to the host served by the storage system called <code>toaster</code>. ◆ <code>sanlun lun show toaster:/vol/vol0/lun0</code> to display a summary listing of all the LUNs available to the host served by <code>lun0</code> on <code>toaster</code>. Note _____ When you specify either the <code>sanlun lun show <storage_syatem_name></code> or the <code>sanlun lun show <storage_syatem_name:storage_syatem_pathname></code> 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. _____

Step	Action																																																																																																																																																												
	Example 2: This example displays sanlun lun show -p output for a system with VERITAS multipathing. <pre>sanlun lun show -p</pre> <pre>dbssystem-04:/vol/vol1/db12_test_0 (LUN 0) DMP_NODE: FAS960_0 3g (3221225472) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C Multipath-provider: Veritas DMP 4.0</pre> <pre>-----</pre> <table><tr><td>host</td><td>filer</td><td></td><td>primary</td><td>partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>device</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>policy</td></tr><tr><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td></tr><tr><td>up</td><td>secondary</td><td>hdisk3</td><td>fcs0</td><td></td><td>9b</td></tr><tr><td>up</td><td>primary</td><td>hdisk8</td><td>fcs1</td><td>5a</td><td></td></tr><tr><td>up</td><td>secondary</td><td>hdisk7</td><td>fcs1</td><td></td><td>5b</td></tr><tr><td>up</td><td>primary</td><td>hdisk4</td><td>fcs0</td><td>9a</td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr></table> <pre>dbssystem-03:/vol/vol1/db12_test_0 (LUN 0) DMP_NODE: FAS960_0 3g (3221225472) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C Multipath-provider: Veritas DMP 4.0</pre> <pre>-----</pre> <table><tr><td>host</td><td>filer</td><td></td><td>primary</td><td>partner</td><td>Veritas</td></tr><tr><td>path</td><td>path</td><td>device</td><td>host</td><td>filer</td><td>filer</td></tr><tr><td>state</td><td>type</td><td>filename</td><td>HBA</td><td>port</td><td>port</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>policy</td></tr><tr><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td><td>-----</td></tr><tr><td>up</td><td>secondary</td><td>hdisk5</td><td>fcs0</td><td></td><td>9b</td></tr><tr><td>up</td><td>secondary</td><td>hdisk9</td><td>fcs1</td><td></td><td>5b</td></tr><tr><td>up</td><td>primary</td><td>hdisk6</td><td>fcs1</td><td>5a</td><td></td></tr><tr><td>up</td><td>primary</td><td>hdisk2</td><td>fcs0</td><td>9a</td><td></td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr><tr><td></td><td></td><td></td><td></td><td></td><td>Balanced</td></tr></table>	host	filer		primary	partner	Veritas	path	path	device	host	filer	filer	state	type	filename	HBA	port	port						policy	-----	-----	-----	-----	-----	-----	up	secondary	hdisk3	fcs0		9b	up	primary	hdisk8	fcs1	5a		up	secondary	hdisk7	fcs1		5b	up	primary	hdisk4	fcs0	9a							Balanced						Balanced						Balanced						Balanced	host	filer		primary	partner	Veritas	path	path	device	host	filer	filer	state	type	filename	HBA	port	port						policy	-----	-----	-----	-----	-----	-----	up	secondary	hdisk5	fcs0		9b	up	secondary	hdisk9	fcs1		5b	up	primary	hdisk6	fcs1	5a		up	primary	hdisk2	fcs0	9a							Balanced						Balanced						Balanced						Balanced
host	filer		primary	partner	Veritas																																																																																																																																																								
path	path	device	host	filer	filer																																																																																																																																																								
state	type	filename	HBA	port	port																																																																																																																																																								
					policy																																																																																																																																																								
-----	-----	-----	-----	-----	-----																																																																																																																																																								
up	secondary	hdisk3	fcs0		9b																																																																																																																																																								
up	primary	hdisk8	fcs1	5a																																																																																																																																																									
up	secondary	hdisk7	fcs1		5b																																																																																																																																																								
up	primary	hdisk4	fcs0	9a																																																																																																																																																									
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								
host	filer		primary	partner	Veritas																																																																																																																																																								
path	path	device	host	filer	filer																																																																																																																																																								
state	type	filename	HBA	port	port																																																																																																																																																								
					policy																																																																																																																																																								
-----	-----	-----	-----	-----	-----																																																																																																																																																								
up	secondary	hdisk5	fcs0		9b																																																																																																																																																								
up	secondary	hdisk9	fcs1		5b																																																																																																																																																								
up	primary	hdisk6	fcs1	5a																																																																																																																																																									
up	primary	hdisk2	fcs0	9a																																																																																																																																																									
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								
					Balanced																																																																																																																																																								

Displaying host HBA information

Displaying host HBA information

To manage the host HBA information, complete the following steps.

Step	Action
1	Log in as root on the host.
2	At the host command line, use the <code>sanlun</code> command to display host HBA information: <code>sanlun fcp show adapter [-c [-v] [adapter name all]]</code> <code>-v</code> produces verbose output. <code>-c</code> gives configuration information that you can use to create igroups. <code>all</code> lists information for all FCP adapters. <code>adapter name</code> lists adapter information for the adapter you specify . Example 1: The following example displays the type of output you see with the <code>sanlun fcp show adapter</code> command line. This example simply displays information on the adapters. The next one includes the <code>-v</code> option, which provides more information on the adapter. <pre># sanlun fcp show adapter fcs0 WWPN:10000000abcde123 fcs1 WWPN:10000000abcde456</pre>

Step	Action
	Example 2: This example adds the <code>-v</code> option to the <code>sanlun fcp show adapter</code> command line. The <code>-v</code> option provides more information on the adapter. <pre># sanlun fcp show adapter -v fcs0 WWPN:10000000abcde123 WWNN: 20000000c936268b driver name: /usr/lib/drivers/pci/efcdd model: df1000f9 model description: FC Adapter serial number: 1E32909972 hardware version: not available driver version: 5.1.0.50 firmware version: CD382A1 Number of ports: 1 port type: Fabric port state: Operational supported speed: 2 GBit/sec negotiated speed: 1 GBit/sec OS device name: fcs0 fcs1 WWPN:10000000c9362687 WWNN: 20000000c9362687 driver name: /usr/lib/drivers/pci/efcdd model: df1000f9 model description: FC Adapter serial number: 1E32909970 hardware version: not available driver version: 5.1.0.50 firmware version: CD382A1 Number of ports: 1 port type: Fabric port state: Operational supported speed: 2 GBit/sec negotiated speed: 1 GBit/sec OS device name: fcs1 #</pre>

About this chapter

This chapter describes how to troubleshoot software installation problems and HBA and LUN management problems.

Topics covered in this chapter

This chapter discusses the following topics:

- ◆ [“Troubleshooting checklists”](#) on page 90
- ◆ [“Installation fails because system does not meet requirements”](#) on page 92
- ◆ [“Commands you can use to troubleshoot problems”](#) on page 93
- ◆ [“SAN Toolkit diagnostic utilities”](#) on page 94
- ◆ [“Change the AIX HBA’s attributes”](#) on page 100
- ◆ [“Change the HBA queue depth with no LUNs mapped”](#) on page 104
- ◆ [“Change the value for the queue depths on the disk”](#) on page 105
- ◆ [“General troubleshooting”](#) on page 106

Troubleshooting checklists

Connectivity problems

Because applications on IBM AIX systems do not typically handle I/O errors gracefully, read/write failures can cause applications to hang, crash, or fail to return requested information.

You can minimize I/O application error on AIX systems by using highly available, multipathing configurations that provide continuous I/O with no single point of failure. If you continue to experience I/O errors, check the following areas for connectivity problems:

- ◆ Storage system setup
- ◆ Storage system connectivity to the Fibre Channel switch
- ◆ Fibre Channel switch setup
- ◆ Fibre Channel connectivity to the host
- ◆ Host setup

Storage setup checklist

The following is a checklist of items that pertain to the setup. You should make sure that the following situations are true:

- ◆ The storage system FCP cfmodes is correctly set for your configuration. See Chapter 4, “[Migrating a Storage Cluster to Single Image Mode](#),” on page 41.
- ◆ FCP is running and all the target ports are in a good state.
- ◆ The igroups have the ostype set to AIX.
- ◆ The LUNs are
 - ❖ mapped to host igroups
 - ❖ in a good state as indicated by the output from the `sanlun lun show -p` command or the `lspath` command.
 - ❖ use AIX as their operating system type

Host setup checklist

The following is a checklist of items that pertain to the host setup. You should make sure that the following is true:

- ◆ The system meets the FCP IBM AIX Host Utilities requirements.
The installation fails if your system does not meet the requirements.
- ◆ The HBAs are working.
Check the output from the `sanlun fcp show adapter -v` command.

- ◆ The host is seeing all of the LUNs and LUN paths that it should.
You can confirm this by running the `lsdev` command or, if you have multipathing, `sanlun lun show -p`.
- ◆ Use VERITAS commands to check your path state and device availability.
- ◆ The `sanlun lun show -p` command is not reporting problems such as
 - ❖ downed paths
 - ❖ missing paths

Installation fails because system does not meet requirements

Installation checks for system requirements

The installation process and requires certain minimum system requirements. If it finds a problem with your system requirements, the installation fails.

Reinstalling the software

If the installation fails because your system does not meet the system requirements, complete the following steps.

Step	Action
1	Check the installation logs to determine why the install failed. The most common problems include missing the required filesets, such as devices.common.IBM.fc.hba-api fileset.
2	Correct the problem described in the log.
3	Run SMIT to install the software. Note _____ Unstalling software is not necessary. _____

Commands you can use to troubleshoot problems

Tools you can use

There are several commands you can use to check out your system and look for problems. The following commands are especially useful in tracking down problems:

- ◆ The SAN Toolkit `sanlun` utility
- ◆ The AIX `lsdev` command
- ◆ The VERITAS commands `vxddmadm` and `vxdisk list`

Using `sanlun`

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

- ◆ Check the status of the
 - ❖ HBA
 - ❖ LUN
 - ❖ Paths to the LUN
- ◆ Discover whether a path is up or down

When you have a direct-attached configuration, use the `sanlun lun show` command to get information on your system.

With a multipathing configuration, use `sanlun lun show -p` command, which includes information on the paths to LUNs.

Examples of the `sanlun` commands are in the section “[Getting the host to discover the new LUNs](#)” on page 64.

Using the `lsdev` command

The `lsdev -C -c disk` command provides information about the devices and LUNs in your configuration. The following example shows the type of output displayed when you run this command. (This example has been shortened.)

```
# lsdev -Ccdisk
hdisk0 Available 08-08-00-5,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 08-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 04-08-02      MPIO NetApp FCP Default PCM Disk
hdisk3 Available 04-08-02      MPIO NetApp FCP Default PCM Disk
hdisk4 Available 04-08-02      MPIO NetApp FCP Default PCM Disk
hdisk5 Available 04-08-02      MPIO NetApp FCP Default PCM Disk
```

SAN Toolkit diagnostic utilities

Diagnostic utilities

The SAN Toolkit provides several utilities that provide diagnostic information you can use for troubleshooting.

When you run these utilities, they gather configuration information about different aspects of your system. They automatically collect information and place it in a directory. Then they create a gzipped .tar file of that directory. You can send this file to technical support.

Normally, you only use these utilities when requested to do so by NetApp’s Customer Support. The utilities are

- ◆ filer_info. See “[Collecting storage system information](#)” on page 95 for more information.
- ◆ aix_info. See “[Using the AIX data-collection tool](#)” on page 96.
- ◆ <switch>_info, where <switch> is the brand of switch you are using (for example, brocade_info). See “[Collecting switch information](#)” on page 97.

General steps for executing the programs

Although some of the scripts have arguments that you may want to use, the general steps for executing these scripts are the same.

Step	Action
1	Go to the directory /opt/NetApp/santools/bin: <code>cd /opt/NetApp/santools/bin</code>
2	Enter the utility name. If the utility you are using includes arguments, you can enter them as well. Result: The configuration information that the utility gathers is placed in a gzipped .tar file.

Collecting storage system information

If you are having a problem, NetApp Customer Support might ask you to run `filer_info` to collect storage system configuration information. You should only run this utility when requested to do so by Customer Support. To execute `filer_info`, complete the following steps.

Note

This utility runs in command-line mode only.

Step	Action
1	Go to the directory <code>/opt/NetApp/santools/bin</code>: <pre>cd /opt/NetApp/santools/bin</pre>
2	At the host command prompt, enter the following command: <pre>filer_info [-d path_and_file_name] [-l user] -n <filer_name></pre> The <code>-d</code> option enables you to specify a different path in which to place the resulting output. The default path is <code>/tmp/NetApp/ntap_filer_<filer_name></code>. The <code>-l</code> user option allows the choice of a user name and password on the switch. The target of the command is specified by <code>-n <filer_name></code>, where <i>filer_name</i> is the name of IP address of the storage system. Result: The output of various status commands are placed in the <code>e</code> output directory named <code>netapp_filer_<filer_name></code>, where <i>filer_name</i> is the storage system name or IP address. The contents of this directory are placed in a tarred and gzipped file named <code>filer_<filer_name>.tar.gz</code>. Example: The following command line checks <code>system_1</code>, logging on to it as the user <code>root</code>. <pre>host1> filer_info -d /opt/NetApp/santools -l root -n system_1</pre> The information from <code>system_1</code> is placed in a gzipped file with the directory path <code>/opt/NetApp/santools/filer_filer1.tar.gz</code> as its name.
3	Send the resulting gzipped file to Customer Support for analysis.

Using the AIX data-collection tool

If you are having a problem, Customer Support might ask you to run `aix_info` to collect configuration information on your AIX system. You should only run this utility when requested to do so by Customer Support.

To execute `aix_info`, complete the following steps:

Step	Action
1	Change to the <code>/opt/NetApp/santools/bin</code> directory: <code>cd /opt/NetApp/santools/bin</code>

Step	Action
2	At the command prompt, enter the following command: <pre>aix_info [-d <directory>] [-n <file_name>]</pre> -d <directory> specifies the location for the gzipped file that aix_info creates. The default location is /tmp/netapp. -n <file_name> specifies the name for the output file and directory names. The default pathname is /tmp/netapp/ntap_aix_info.tar.gz. Note Ensure that you have at least 64MB of free storage in the output directory Result: When you run aix_info, the utility gathers information about your AIX configuration and places it in a tarred, gzipped file. It also creates a directory and populates it with files containing output from numerous commands run on the AIX system and copies of logs such as the output of the errpt command. Example: Depending on the options you specify, you see output similar to the following when you run this utility: <pre># aix_info # aix_info Collecting aix_info.....(collecting AIX snap... may take several minutes)..... AIX system info is in directory /tmp/netapp/ntap_aix_info Compressed file is /tmp/NetApp/ntap_aix_info.tar.Z Please send this file to Network Appliance for analysis</pre>
3	Send the resulting gzipped file to Customer Support for analysis.

Collecting switch information

If you are having a problem, Customer Support might ask you to run the switch utility the SAN Toolkit supplies when you install the toolkit software. The SAN Toolkit provides a utility for each switch it supports for this Host Utility. The

utility name is <switch>_info, where *switch* is the brand of the switch (for example, if you have a McDATA switch, the file name would be mcddata_info). This utility collects switch configuration information. You should only run this utility when requested to do so by Customer Support.

To execute the utility, complete the following steps.

Note

This utility runs in command-line mode only.

Step	Action
1	Go to the directory /opt/NetApp/santools/bin: <code>cd /opt/NetApp/santools/bin</code>

Step	Action
2	At the host command prompt, enter the following command: <pre><switch_brand>_info [-d path_and_file_name] [-l user:password] -n <switchname></pre> <switch_brand> is the name of the switch you are using; for example, brocade for a Brocade switch. The -d <i>path_and_file_name</i> option enables you to specify a different path in which to place the resulting output. The default path is /tmp/netapp/<switch_brand>_switch_<switch_name>. The -l <i>user:password</i> option allows the choice of a user name and password on the switch. If this option is not specified, the switch default login and password settings are used. The target of the command is specified by -n <i>switchname</i>, where <i>switchname</i> is the name or IP address of the switch. Result: The output of various status commands on the switch are placed in a directory named <switch_brand>_switch_<switchname>, where <i>switch_brand</i> is the brand of the switch, such as Brocade, and <i>switchname</i> is replaced by the switch name or IP address. The contents of this directory are placed in a tarred and gzipped file named <switch_brand>_switch_<switchname>.tar.gz. Example: This example uses the brocade_info file. The switch brcd1 is connected with the user "john", who has the password "smith". <pre>host1> brocade_info -d /opt/netapp/santools -l john:smith -n brcd1</pre> The information is placed in a tarred, gzipped file that has the pathname /tmp/netapp/brocade_switch_brcd1.tar.gz.
3	Send the resulting gzipped file to Customer Support for analysis.

Change the AIX HBA’s attributes

Methods for changing the FCP HBA’s attributes

If you want to change information, such as the queue depth on the HBA, you must change the HBA’s attributes. If you have LUNs mapped to the host, there are two methods you can use to change these attributes. One method requires you to reboot the host for the change to take effect. The other method does not require a reboot, but you must temporarily unload all the HBA’s child devices.

Note—————
You must unload the child devices even if you are using SMIT. SMIT does not unload these devices automatically.

If you don’t have any LUNs mapped to the host, you can use a shorter way to change such attributes as the HBA queue depth (see “[Change the HBA queue depth with no LUNs mapped](#)” on page 104).

Changing the FCP HBA’s attributes and rebooting

If you only change the FCP HBA’s attributes in the ODM, you must reboot your host. To change the attributes this way, complete the following steps.

Step	Action
1	Enter the following command: <pre># chdev -P -l <i>adapter_name</i> -a num_cmd_elems=<i>val</i></pre> <i>adapter_name</i> is the name of the HBA whose values you want to change. <i>val</i> is the value to which you want to set the queue depth. Example: The following example sets the queue depth on the FCP HBA device named fcs1 to 400. <pre># chdev -P -l fcs1 -a num_cmd_elems=400</pre>
2	Reboot the host.

Changing the FCP HBA's attributes without rebooting

If you don't want to reboot the host, you can change the HBA attributes by temporarily unloading all the HBA child devices. To use this method to change the attributes, complete the following steps.

Note

The commands in these steps provide an example of changing the attribute. Depending on how your system is set up, you would probably see different output if you executed these commands on your system.

Step	Action
1	Use the <code>lsdev</code> command to find the HBA's child device. This is a device instance of SCSI over the FCP driver. If this command does not return any child devices, go to step 11. In this case, you will stop at step 12. Example: In this example, the <code>lsdev</code> command reports one FCP device. <pre># lsdev -C -F "name parent" awk '\$2~/^fcs1\$/ { print \$1 }'</pre> fscsi1
2	Now that you have a name for the FCP device, use the <code>lsdev</code> command to locate that device's child devices. These devices are among the LUN disk devices. If this command doesn't display any child devices, go to step 10 and stop after step 13. Example: In this example, the <code>lsdev</code> command displays four child devices. <pre># lsdev -C -F "name parent" awk '\$2~/^fscsi1\$/ { print \$1 }'</pre> hdisk2 hdisk3 hdisk6 hdisk7
3	Use the <code>lspv</code> command to find all the volume groups that contain at least one disk device that is a child of the FCP device. If the command does not report any volume groups, skip to step 8 and stop after step 14.

Step	Action
4	Find all the mounted file systems that belong to those volume groups and the applications that are currently accessing logical volumes in those volume groups. If there aren't any, skip to step 7 and stop after step 15.
5	Disable any applications that are accessing logical volumes in the volume groups.
6	Vary off all the volume groups and unmount any file systems. Example: To vary off the volume group called my_vg, you would enter the following command. <pre># varyoffvg my_vg</pre>
7	Unload all the protocol's child devices. You need to keep the definition of these devices. Example: This example unloads the child devices that were discovered in step 2. <pre># rmdev -l hdisk2 hdisk2 Defined # rmdev -l hdisk3 hdisk3 Defined # rmdev -l hdisk6 hdisk6 Defined # rmdev -l hdisk7 hdisk7 Defined</pre>
8	Unload the FCP device that you discovered in step 1. Example: This example unloads the FCP device fscsi1. <pre># rmdev -l fscsi1 fscsi1 Defined</pre>
9	Change the value of the attribute. Example: This example changes the queue depth value to 400. <pre># chdev -l fcs1 -a num_cmd_elems=400 fcs1 changed</pre>

Step	Action
10	Reload the HBA device to enable the new setting. Example: This example reloads the FCP device fcs1. <pre># mkdev -l fcs1 fcs1 Available</pre>
11	If necessary, reload the FCP device fcs1. Example: The following command line reloads the FCP device fcs1. <pre># mkdev -l fcs1 fcs1 Available</pre>
12	If necessary, reload the child devices of the FCP device. Example: The following command line reloads the child device that you unloaded in step 9. <pre># mkdev -l hdisk2 hdisk2 Available # mkdev -l hdisk3 hdisk3 Available # mkdev -l hdisk6 hdisk6 Available # mkdev -l hdisk7 hdisk7 Available</pre>
13	Vary on any volume groups that you varied off in step 7. Example: To vary on the volume group called my_vg, you would enter the following command. <pre># varyonvg my_vg</pre>
14	Remount any file systems that you unmounted.
15	Restart any applications that you halted.

Change the HBA queue depth with no LUNs mapped

Changing the HBA queue depth

If you need to change the HBA queue depth, and you don't have any LUNs mapped to the host, complete the following steps.

Note—————
If you have LUNs mapped to the host, you must perform the steps listed in [“Changing the FCP HBA’s attributes without rebooting”](#) on page 101.

Step	Action
1	Unload the child FCP device by using the <code>rmdev -l</code> command. Example: For HBA <code>fscsi0</code> , you would enter the following command: <code># rmdev -l fscsi0</code>
2	Change the value of the queue depth using the <code>chdev -l</code> command, which has the following format: <code># chdev -l fcs0 -a num_cmd_elems=val</code> <i>val</i> is the value to which you want to set the queue depth. Example: To change the value to 48, you would enter the following command. <code># chdev -l fcs0 -a num_cmd_elems=48</code>
3	Reload the HBA device using the <code>mkdev -l</code> command. Example: The following command line reloads the <code>fcs0</code> HBA. <code># mkdev -l fcs0</code>
4	Reload the child FCP device using the <code>mkdev -l</code> command. Example: The following command line reloads the child FCP device. <code># mkdev -l fscsi0</code>

Change the value for the queue depths on the disk

Setting the value of the queue depths

The installation automatically sets the queue depths for the disk to 12, which is the recommended value. If your system needs to use a different queue depth, change this value using the `chdev -l` command. This command has the following format:

```
# chdev -l hdiskx -a queue_depth=val
```

`hdiskx` is the device for which you want to change the depth. The *x* specifies the number for this device.

val is the value to which you want to set the queue depth.

Note

You should only change the queue depth if you are familiar with performance tuning on your system.

Example of changing the queue depth

This example sets the queue depth for the `hdisk6` to 12.

```
# chdev -l hdisk6 -a queue_depth=12
```

General troubleshooting

Basic operations take a long time

Problem: There are no errors and everything appears to be correctly connected, but basic operations are taking exceptionally long times; in some cases, hours instead of minutes.

Explanation: If there are no other problems, this could mean that one of the connectors is not firmly attached. This can be the case even when it has a link and is reporting information such as the correct speed.

What to do: Disconnect the connector and reconnect it. Check for faulty optical cabling. Also, check the fibre channel switch for excessive line errors.

Index

A

AIX host

- example of successful installation 24
- example of successful uninstallation 34
- uninstalling Host Utilities software 29
- using SMIT 19

Array Support Library

- see ASL 36

ASL

- definition 36

attributes

- changing and rebooting 100
- changing on HBA 100
- changing with SMIT 100
- changing without rebooting 101

C

CD-ROM

- getting software 16

cfgmgr command 64, 78

cfmode

- correct setting 42

chdev command 100

checklists

- filer setup 90
- host setup 90
- installation tasks 2
- troubleshooting 90

configuration

- direct-attached 56

D

diagnostic utilities

- aix_info 94
- executing 94
- filer_info 94, 95
- switch_info 94, 97

direct-attached configuration

- set media type 56

disk

- changing value of queue depths 105

disk shelves

- powering up 10

E

equipment

- powering up 10

F

file system

- creating 67, 73

filer_info

- executing 95

files

- filer_info 95
- switch_info 97

fileset

- check version 7

H

HBA

- changing attributes 100
- changing queue depth 104
- installing 8
- verifying queue depth 28

host 14

- check OS version 7
- communicating with storage system 10
- creating volume group 65, 71
- device for physical volume 70
- discovering new LUNs 64, 78
- displaying LUN information 84
- powering up 10
- requirements 6
- setup checklist 90

I

- igroup show command 59
- install command
 - checks system requirements 92
 - rerunning 92
- installation
 - pre-installation tasks 6

L

- lsdev command 101
 - using for troubleshooting 93
- lslpp command 7
- lun commands
 - lun create 60
 - lun map 60
- LUN ID
 - defined vii
- lun setup command 59
- lun show 61
- LUNs
 - best practices for creating 57
 - creating with lun setup 59
 - defined vii
 - discovering 64, 78
 - displaying host information 84
 - tasks for creating on filer 57
 - verifying 61

M

- media type
 - setting 56

N

- NOW site
 - downloading software 15

O

- oslevel command 7

P

- physical volume
 - selecting device 70
- powering up equipment 10

Q

- queue depth
 - changing value on disk 105
 - changing value on HBA 104
 - verifying 28

R

- requirements 6

S

- sanlun fcp show adapter command 58, 87
- sanlun lun show command 84
- sanlun utility
 - about 83
 - sanlun fcp show adapter 87
 - using for troubleshooting 93
- SMIT
 - creating file system 67, 73
 - creating volume group 65, 71
 - unloading child devices 100
 - using on AIX 19
- software
 - downloading 15
 - getting 14
 - getting from CD-ROM 16
 - installation tasks 18
 - uncompressing 16
- storage
 - accessing on volume group 67, 73
- storage system
 - communicating with host 10
 - correct dfmode 42
 - direct-attached 56

- powering up 10
- setup checklist 90
- verifying setup 8
- switch_info
 - executing 97
- switches
 - powering up 10
 - verifying setup 9
 - zoning 9
- system requirements 6

T

- Terminology vii
- troubleshooting
 - checklists 90
 - using lsdev command 93

- using sanlun 93

U

- uninstalling
 - Host Utilities software 29

V

- volume group
 - accessing storage 67, 73
 - creating 65, 71

Z

- zoning 9

