

FCP IBM® AIX® Attach Kit 1.3

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 215-01767_A0
May 2006

Copyright and trademark information

Copyright information

Copyright © 1994—2005 Network Appliance, Inc. All rights reserved.

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, Data ONTAP, FilerView, MultiStore, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. Network Appliance, and Snapshot are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. FlexClone, FlexVol, NOW, NOW NetApp on the Web, RAID-DP, SecureAdmin, SnapDrive, vFiler, VFM, and Virtual File Manager, are trademarks of Network Appliance, Inc. in the United States and other countries.

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	v
Chapter 1	Installation and Setup Overview	1
	FCP IBM AIX Attach Kit 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.	22
Chapter 3	Installing the ASL and APM	27
	Overview of the Array Support Library and Array Policy Module	28
	Installing the ASL and APM	29
Chapter 4	Migrating a Storage Cluster to Single Image Mode.	33
	Understanding cfmode	34
	Tasks to complete before you begin	38
	Upgrading to single image mode	40
Chapter 5	Creating, Configuring, and Managing LUNs	47
	Preparing the storage system	48
	Creating and mapping igroups and LUNs	49
	LUN configuration overview.	54
	Configuring LUNs with native AIX LVM.	56
	Configuring LUNs with SANpath multipathing software	63
	Configuring LUNs with VxVM and DMP multipathing software.	73
	Displaying LUN and HBA information	80
	Using sanlun to display host LUN information	81

	Displaying host HBA information	85
Chapter 6	Configuring SAN booting.	87
	SAN boot overview.	88
	Creating the boot LUN.	90
	Booting the AIX Installation CD-ROM	94
	Installing the OS	95
	Troubleshooting the SAN boot LUN.	102
Chapter 7	Troubleshooting	103
	Troubleshooting checklists	104
	Installation fails because system does not meet requirements	106
	Commands you can use to troubleshoot problems	110
	Use SANpath commands when working with multipathing	112
	Attach Kit diagnostic utilities.	114
	Change the AIX HBA's attributes	119
	Change the HBA queue depth with no LUNs mapped	124
	Change the value for the queue depths on the disk.	125
	General troubleshooting	126
	Index	127

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:

- ◆ *igroup* 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.
- ◆ *LUN* (Logical Unit Number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.
- ◆ *Initiator* refers to the system component that originates an I/O command over an I/O bus or network.
- ◆ *Target* refers to the storage system that receives an FCP I/O request and serves the data requested by the initiator.
- ◆ *Type* refers to pressing one or more keys on the keyboard.
- ◆ *Enter* 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 Attach Kit.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“FCP IBM AIX Attach Kit 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 Attach Kit installation checklist

List of tasks to perform

Installing the FCP IBM AIX Attach Kit 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 Attach Kit does not include the host HBA.	Host	“ Verifying 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> ◆ <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	“Where to go for more information” on page 11
Installing the software package			
8	Get a copy of the Attach Kit software package.	Host	“Ways to obtain the software packages” on page 14
9	If you downloaded the software package, uncompress it and extract the files. If you are using the FCP IBM AIX Attach Kit 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 software package. If your configuration requires multipathing, install one of the following: ◆ Dot Hill SANpath package — available on the NetApp NOW™ site and on the CD-ROM (SANpath version only) ◆ 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 ◆ The FCP Host Compatibility Matrix at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml for a direct link to the ASL and APM on the Symantec Web site. ◆ “Configuring LUNs with native AIX LVM” on page 56 ◆ “Configuring LUNs with SANpath multipathing software” on page 63 ◆ “Configuring LUNs with VxVM and DMP multipathing software” on page 73 ◆ “Installing the ASL and APM” on page 27
11	Make sure the HBA initiator queue depth is appropriate for your system.	Host	“ Verifying the HBA initiator queue depth ” on page 21
Creating and configuring LUNs			
12	If you have a storage cluster, it is recommended to migrate your cfmodesetting 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 33 ◆ “Setting the media type” on page 48

Step	Action	Machine	Details
13	Create igroup(s) using AIX as the value for the <code>ostype</code> attribute.	Host, Storage System	“Creating an igroup” on page 49
14	Create LUN(s) that map to the igroup.	Storage System	“Creating and mapping LUNs” on page 51
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 56
16	Configure the LUNs ◆ If your configuration does not require multipathing, use native AIX LVM to configure LUNs. ◆ If your configuration requires multipathing, use native AIX with SANpath or VERITAS software to configure LUNs.	Host	Refer to: ◆ “Configuring LUNs with native AIX LVM” on page 56 ◆ “Configuring LUNs with SANpath multipathing software” on page 63 ◆ “Configuring LUNs with VxVM and DMP multipathing software” on page 73.
17	Display information about the LUNs and HBA.	Host	Refer to: ◆ “Using sanlun to display host LUN information” on page 81 ◆ “Displaying host HBA information” on page 85

Completing the pre-installation tasks

Tasks to perform

Before you install Attach Kit software, verify the following:

- ◆ Host system requirements
- ◆ Storage system installation and configuration
- ◆ Switch installation and configuration
- ◆ Storage/host and switch cabling
- ◆ Host and storage system communication

Verifying that the host meets 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>.

Dynamic LPAR

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.

Note

SANpath does not support DLPAR with HBAs.

Check your operating system version

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

```
oslevel -r
```

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

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

```
# oslevel -r
5200-04
```

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
R -- RPM Package
```

Verifying the host HBA

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

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 have not installed the 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 Attach Kit, and to 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.

Determine 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.

If you have a high-availability configuration, make sure that there is at least one primary path and one secondary path through each switch for each LUN. 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>.

Note

The SANpath multipathing software limits the number of paths per LUN to a maximum of 32. If your host has more than 32 paths, you can use switch zoning to reduce the number of those paths that are visible to the LUN.

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

Where to go for more information

Host and storage system information

Use the documents below to access 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...
Multipathing	◆ <i>SANpath User's Guide for IBM AIX Operating Systems</i> — describes how to use SANpath software to configure multipathing on a storage device. ◆ Chapter 3, “Installing the ASL and APM,” on page 27 ◆ <i>FCP Quick Reference for FCP IBM AIX Attach Kit 1.3</i> — provides command syntax for native AIX LVM configurations, and configurations with SANpath or VERITAS multipathing.
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	◆ <i>Data ONTAP™ Block Access Management Guide</i> ◆ <i>Data ONTAP Release Notes</i>
Verifying compatibility of storage systems with environmental requirements	◆ <i>Site Requirements Guide</i>
Upgrading Data ONTAP	◆ <i>Data ONTAP Upgrade Guide</i>

About this chapter

This chapter describes procedures for installing the FCP IBM AIX Attach Kit 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 22

Obtaining the software packages

Software packages

The FCP IBM AIX Attach Kit includes:

- ◆ The FCP IBM AIX Attach Kit software— contains NetApp's FCP device definitions that are required for FCP IBM AIX configurations.
- ◆ The NetApp toolkit software- 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.

Configurations that include multipathing require one of the following:

- ◆ Dot Hill SANpath multipathing software package, available with the SANpath version of the kit.
- ◆ 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 Attach Kit on the NOW site provides the Attach Kit software packages bundled into a single compressed .tar file. SANpath multipathing software is available from the same site but it requires an additional license. Detailed information about obtaining the SANpath Activation Key are available on the download page for the Attach Kit software. For information about downloading this file, see “[Downloading the software](#)” on page 15.
- ◆ Get them from the CD-ROMs provided with the media package in your FCP IBM AIX Attach Kit. 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 (FCP) Host Attach Kit 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 software file to a local directory.
6	To get the software necessary for multipathing, download the SANpath software file to the same directory where you downloaded the NetApp software.
7	If you download the files to a machine other than the AIX host, you'll need to move them to the host.

Uncompressing the software file you downloaded

To uncompress the software file 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 command to uncompress the FCP IBM AIX Attach Kit package and extract the software: <pre># zcat ntap_aix_fcp_1.3.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 produces output similar to the following. <pre># zcat ntap_aix_fcp_1.3.tar.Z tar xvf - x ./ntap_aix_fcp_1.3/NetApp.fcp_attach_kit, 22528 bytes, 44 media blocks. x ./ntap_aix_fcp_1.3/NetApp.fcp_toolkit, 1005568 bytes, 1964 media blocks. x ./ntap_aix_fcp_1.3/install, 37326 bytes, 73 media blocks. x ./ntap_aix_fcp_1.3/uninstall, 17844 bytes, 35 media blocks. #</pre>
2	If you are using SANpath multipathing software, enter the following command to uncompress and extract the software. <pre># zcat dothill_sanpath_4.0.5n.tar.Z tar -xvf -</pre>
3	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. Enter <code>cd <CDROM mountpoint></code>

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

Installing the software packages

Tasks involved in installing the Attach Kit software

The way you use the `install` command depends on your configuration.

- ◆ SANpath configurations—Enter the `install` command without arguments.
- ◆ VERITAS—Enter the `install` command with the `-n` argument.

What the install command does

When you run the installation command, it performs the following tasks:

- ◆ Displays the operating system version, HBA model, and firmware of your host.
- ◆ Installs the toolkit package.
- ◆ Depending on the installation option you select, changes the host HBA ODM values to interoperate with NetApp storage systems.
- ◆ Runs the `bosboot` command to update the boot image.
- ◆ Logs all install command output to a file named `netapp_install_log_yymmddhhMMSS` where `yymmddhhMMSS` represents the date and time that you started the installation. By default this log file is placed in `/tmp` and `/opt/NetApp`.

Caution

The FCP IBM AIX installation script does not prevent you from installing unsupported components. To ensure your configuration meets requirements, check the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_icsi_config/.

Upgrade

It is recommended to remove your current installation of the attach kit software before you install a new package. See “[Uninstalling the software packages](#)” on page 22.

Installing the Attach Kit To install the Attach Kit software for SANpath configurations, complete the following steps.

Step	Action
1	Log in as root and change to the directory that contains the FCP IBM AIX Attach Kit software. SANpath configurations—Enter the following command <pre>./install</pre> VERITAS configurations or configurations that do not have multipathing—Enter the following command: <pre>./install -n</pre> Result: The install command checks and displays host configuration information. <pre>NetApp Attach Kit for AIX install</pre> <pre>This host is running AIX 5.3 ML2. There are 2 HBAs detected on this host. fcs0 is an FC6239. The firmware (microcode) level is 1.90X2. fcs1 is an FC6239. The firmware (microcode) level is 1.90X2. Please check the NetApp SAN Support Matrix for OS and HBA Compatibility. OS and HBA check complete. WARNING: SANpath will not be installed. Please consult your manual to understand how this impacts the availability of your NetApp filer to this host.</pre> Note_____ Disregard the warning that states SANpath will not be installed. The install command does not automatically install any multipathing software. If you have SANpath software, you must install it separately. _____
2	Specify the location of the software package. Result: The installation command prompts you for the location of the packages, and installs the software. Please enter the location of the NetApp packages: (if you are installing from CD, you may enter the device name, e.g. /dev/cd0) [default: /tmp/1.3/ntap_aix_fcp_1.3]

Step	Action
3	The installation command asks how you want to handle updating the Object Data Manager (ODM) settings for each HBA in your system. Note Unless you have specific requirements for tuning your system, it is recommended that you choose option 3, which updates the settings and saves the old HBA settings to the file <code>/opt/NetApp/fcp_attach_kit/saved_hba_settings</code>. Option 3 enables you to update the settings for each HBA connected to the storage system rather than updating the settings for all host HBAs. The NetApp AIX FCP Attach Kit 1.3 requires changes to device attributes for the Fibre Channel adapter <code>fcs0</code>. This script can update the ODM settings, which will take effect at the next reboot. You can either: 1) Update the ODM settings for all HBAs; 2) Skip update for all HBAs. 3) Update the ODM settings for this HBA; 4) Do not update anything for this HBA, skip to the next one; What would you like to do? If you are not certain, NetApp recommends selecting the default [1]: <pre>> 3</pre> The NetApp AIX FCP Attach Kit 1.3 requires changes to device attributes for the Fibre Channel adapter <code>fcs1</code>. This script can update the ODM settings, which will take effect at the next reboot. You can either: 1) Update the ODM settings for all HBAs; 2) Skip update for all HBAs. 3) Update the ODM settings for this HBA; 4) Do not update anything for this HBA, skip to the next one; What would you like to do? If you are not certain, NetApp recommends selecting the default [1]: <pre>> 3</pre> HBA settings update successful. Running <code>bosboot</code> to update the boot image... Installation of the NetApp AIX Attach Kit was successful. You MUST reboot the system before accessing your NetApp filer(s). Installation log saved in both <code>/tmp/netapp_install_log_060516104006</code> and <code>/opt/NetApp/install_log</code> Result: The installation command completes by running <code>bosboot</code> to rebuild the boot image. The changes to the ODM settings do not take effect until you reboot the system.

Step	Action
4	It is recommended to reboot the system when the installation completes. <code>shutdown -F -r now</code>

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 Attach Kit 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 40.

Methods of uninstalling the software

You can `uninstall` the software by running the `uninstall` command interactively, which prompts you for information, or non-interactively, by typing the command and options on the command line.

Note

By default, the `uninstall` command is placed in the directory `/opt/NetApp/santools/bin` when you install the software package.

What the `uninstall` command does

The `uninstall` command performs the following tasks:

- ◆ Removes NetApp devices
- ◆ Optionally restores HBA settings it saved during the installation from the `/opt/NetApp/fcp_attach_kit/saved_hba_settings` file.
- ◆ Uninstalls the toolkit packages.
- ◆ Runs `bosboot` to update the boot image.
- ◆ Logs all command output to a file named `netapp_uninstall_log_yymmddhhMMSS` where `yymmddhhMMSS` represents the date and time that you started uninstalling. By default, the log is saved in the `/tmp` directory.

Uninstalling using the command interactively

To uninstall the attach kit software using the command in interactive mode, complete the following steps

Step	Action
1	Ensure that you are logged in as root.
2	Go to the directory where you uncompressed the software file. For example, if you extracted the software to /tmp, you would enter <code>cd /tmp</code>
3	Use the <code>uninstall</code> command to remove the software. Answer <code>y</code> (or <code>yes</code>) when the command asks if you are sure you want to remove the attach kit. If you answer <code>n</code> (or <code>no</code>), the command does not uninstall the attach kit. <code># ./uninstall</code> This program will uninstall the NetApp Attach Kit. It will attempt to remove NetApp hdisk device entries from the host, and optionally uninstall any NetApp installed multipathing software, and/or restore your pre-installation HBA settings, prompting you if needed (unless you select non-interactive mode). It requires that this host have no volume groups/disk groups active on the NetApp devices. Do you want to continue? [no]yes

Step	Action
4	Specify whether you want to restore your preinstallation HBA settings. ◆ n, which is the default, tells the command to keep the current settings. ◆ y tells the command to restore the previous settings from the file it created during the installation. You must reboot the host after the <code>uninstall</code> command finishes for these settings to take effect. This script can restore your saved pre-installation HBA settings. Would you like to do that? [no] y Restoring pre-installation HBA settings... done.
5	If you are using SANpath multipathing software, specify whether you want to remove or save it. The script retains or removes the multipathing software and removes the packages. SANpath example: The script displays the following if SANpath software is detected: Would you like to uninstall SANpath from your system? [yes] y Uninstalling SANpath... done. Removing NetApp packages...done:
6	If you choose to restore the previous HBA settings, you must enable the restoration by rebooting the host once the <code>uninstall</code> completes. Otherwise, it will continue to use the settings specified during the installation. Running bosboot to update the boot image... done. Uninstallation of the NetApp AIX Attach Kit was successful. You must reboot the system to finish restoring your previous HBA settings. Uninstall log saved in /tmp/netapp_uninstall_log_030930122615

Step	Action
7	If you want to install the latest software version, go to “Uncompressing the software file you downloaded” on page 16.

Using the uninstall command with options

You can uninstall the FCP IBM AIX Attach Kit, the multipathing software, and configure your system by entering the `uninstall` command with the appropriate options. To use this command, you must not have any volume groups defined on your devices. You can use the `uninstall` command in command line mode to clean up after a failed installation.

Step	Action	
1	Log in as root.	
2	Enter the appropriate options on the <code>uninstall</code> command line. <code>uninstall [-f] [-n] [-r]</code>	
	Enter...	To specif ...
	<code>-f</code>	Command line mode, not interactive mode. By default this option uninstalls SANpath and does not restore pre-installation HBA settings.
	<code>-n</code>	That you want to keep the SANPath multipathing software installed on your system.
	<code>-r</code>	That you want the command to restore the HBA settings it saved during the installation.
3	When the <code>uninstall</code> command finishes, reboot the system.	

Example of output using uninstall non-interactively

If you do not want to be prompted for information, you can execute `uninstall` with command line options. When you use the `uninstall` command non-interactively, the command preforms the tasks specified by the options.

This example uses only the `-f` option, which tells the command not to run interactively. Because no other options are included on the command line, the command does not restore the previous HBA settings, but it does remove SANpath (if it has been installed).

```
# /opt/NetApp/santools/bin/uninstall -f
HBA settings not restored.
Removing NetApp packages... done.
Uninstalling SANpath... done.
Running bosboot to update the boot image... done.
Uninstallation of the NetApp AIX Attach Kit was successful.
Uninstall log saved in /tmp/netapp_uninstall_log_03120816172
```

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.

If you are not using VERITAS Storage Foundation 4.0, you do not need the ASL and APM. Proceed to Chapter 5, “[Creating, Configuring, and Managing LUNs](#),” on page 47.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Overview of the Array Support Library and Array Policy Module](#)” on page 28
- ◆ “[Installing the ASL and APM](#)” on page 29

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 33.
- ◆ 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 49.
- ◆ On the host, discover the new LUNs and configure them to be managed by VxVM. See “[VERITAS software \(multipathing\)](#)” on page 55.

About this chapter

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

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' `cfmode` setting to single image.

When you upgrade to Data ONTAP 7.1 from a previous release, your `cfmode` settings are preserved. Your `cfmode` 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 `cfmode` settings, enter the following command on each system in the cluster:

```
fcpx show cfmode
```

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 34.

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 `cfmode` settings on your cluster and reconfigure your host.
- ◆ If you want to keep the current `cfmode` setting for your cluster, proceed to Chapter 5, [“Creating, Configuring, and Managing LUNs,”](#) on page 47.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Understanding cfmode”](#) on page 34
- ◆ [“Tasks to complete before you begin”](#) on page 38
- ◆ [“Upgrading to single image mode”](#) on page 40

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 cfmode setting.
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/<i>filesystem_path</i>
5	Quiesce the volume groups by entering the following command: varyoffvg <i>volume_group_name</i>
6	Export the volume group by entering the following command: exportvg <i>volume_group_name</i>
7	Stop SANpath software by entering the following command: setsp -T
8	Verify that the devices are unavailable by entering the following command: setsp -v If the devices are not available, the <code>setsp</code> command returns a message indicating that no spd devices are found.
9	Enter the following command to list local devices: lsdev -Cc disk
10	Remove the existing NetApp LUN handles for the devices by entering the following command for each device: rmdev -dl <i>device_name</i> Example: <code>rmdev -dl hdisk2</code>

Step	Action
11	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>
12	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 43.
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 and SANpath restarts. If you created volume groups with the autovaryon flag, then the volume groups are reimported.

Step	Action																																																		
3	Verify the new paths by entering the following command. <pre>setsp -v</pre> Result: The <code>setsp -v</code> command displays all available paths to the LUN with a disk handle. The disk handle is the local device SANpath uses as the reference point for all paths to a LUN. The command also activates your volume groups. Example: The example below shows four paths to a LUN. The output <code>spd0=hdisk2</code> indicates that SANpath uses <code>hdisk2</code> to access the LUN. <pre>host1 > setsp -a</pre> <pre>-----</pre> <table><thead><tr><th>spd</th><th>Path/disk</th><th>Status</th><th>Pri</th><th>Exc</th><th>Buf</th><th>Balance</th><th>RtrCnt</th><th>RtrDly</th><th>FailBack</th></tr></thead><tbody><tr><td>0</td><td>hdisk3 (24,1)</td><td>P Good</td><td>X</td><td></td><td>32</td><td>1</td><td>9</td><td>1000</td><td>1</td></tr><tr><td></td><td>hdisk5 (24,3)</td><td>P Good</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>hdisk2 (24,2)</td><td>S Good</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td></td><td>hdisk4 (24,0)</td><td>S Good</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr></tbody></table> <pre> spd0 = hdisk2 ID = "NETAPPData ONTAP LUN OdD0/YnXKxgs" -----</pre>	spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack	0	hdisk3 (24,1)	P Good	X		32	1	9	1000	1		hdisk5 (24,3)	P Good									hdisk2 (24,2)	S Good									hdisk4 (24,0)	S Good							
spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack																																										
0	hdisk3 (24,1)	P Good	X		32	1	9	1000	1																																										
	hdisk5 (24,3)	P Good																																																	
	hdisk2 (24,2)	S Good																																																	
	hdisk4 (24,0)	S Good																																																	
4	If you did not create the volume groups with the <code>autovaryon</code> flag, enter the following command: <pre>importvg -y volume_group_name disk_handle</pre> <i>volume_group_name</i> is the name you assigned to the volume group. <i>disk_handle</i> is the local device SANpath uses as the reference point for all paths to a LUN. The <i>disk_handle</i> is displayed in the output of the <code>setsp -a</code> command you entered in the previous step.																																																		
5	Enter the following command to verify the paths to the LUNs. <pre>sanlun lun show -p</pre>																																																		
6	Verify that the volume groups are active by entering the following command. <pre>lsvg</pre>																																																		
7	If you created file systems on the devices, mount the filesystems.																																																		

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 48
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 49
- ◆ [“LUN configuration overview”](#) on page 54
- ◆ [“Configuring LUNs with native AIX LVM”](#) on page 56
- ◆ [“Configuring LUNs with SANpath multipathing software”](#) on page 63
- ◆ [“Configuring LUNs with VxVM and DMP multipathing software”](#) on page 73
- ◆ [“Displaying LUN and HBA information”](#) on page 80

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 *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 103 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 49. 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 size -t ostype lun-path</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 lun-path igroup-name LUN-ID</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 49. <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 lun-path</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 56
- ◆ If your configuration includes multipathing with SANpath software, refer to [“Configuring LUNs with SANpath multipathing software”](#) on page 63
- ◆ If your configuration includes multipathing with VERITAS software, refer to [“Configuring LUNs with VxVM and DMP multipathing software”](#) on page 73
- ◆ If you are unsure of your configuration and software, refer to [“LUN configuration overview”](#) on page 54

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:

- ◆ Native AIX LVM with SANpath software
- ◆ 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 27.

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 56 for additional information.

SANpath software (multipathing)

To configure LUNs using native AIX LVM with SANpath:

1. Use `cfgmgr` to get the host to discover the new LUNs.
2. Run `sanlun` or SANpath software to display information that helps you determine which device to use as the physical volume.
3. Create a volume using information gathered in the previous step.
4. Configure data access (raw, jfs, or jfs2 file system)

Refer to “[Configuring LUNs with SANpath multipathing software](#)” on page 63 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 `vxdtl 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 73 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 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 colspan="3">F8=Image</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 colspan="4">n=Find Next</td></tr><tr><td colspan="4"><F10 exits smit></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 57. ----- 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									

Configuring LUNs with SANpath multipathing software

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: <pre># cfgmgr</pre> 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. When the device entries are created, SANpath software performs the following tasks: ◆ Detects the new <code>hdisk</code> paths to a LUN. ◆ Creates SANpath multipathing control devices for each LUN. ◆ Starts multipathing control for all <code>hdisk</code> paths to a LUN.
3	Enter <code>sanlun lun show -p</code> to verify that the host has discovered the LUNs. In addition to displaying the LUNs, the output from <code>sanlun lun show -p</code> tells you which paths are primary and which are secondary. Write down which paths SANpath indicates as the access device (<code>SPDn=hdiskn</code>), so that you will have a copy of that information when you create volume paths. (You can also use the SANpath to get information about the host paths. See “Use SANpath commands when working with multipathing” on page 112 for more information.)

Step	Action
	Example: This example shows the type of output you would see in a multipathing environment with SANpath. (This example has been shortened.) <pre> sanlum lun show -p target1:/vol/vol1/db-lun4 (LUN 4) spd8 = hdisk6 100m (104857600) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C Multipath-provider: Dot Hill SANpath version 4.0.5n ----- host filer primary partner SANpath path path device host filer filer failover state type filename HBA port port priority ----- up secondary hdisk31 fcs1 5b failover up primary hdisk26 fcs1 5a primary up secondary hdisk21 fcs0 9b failover up primary hdisk6 fcs0 9a primary target2:/vol/vol1/db-lun3 (LUN 3) spd6 = hdisk5 100m (104857600) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C Multipath-provider: Dot Hill SANpath version 4.0.5n ----- host filer primary partner SANpath path path device host filer filer failover state type filename HBA port port priority ----- up secondary hdisk30 fcs1 5b failover up primary hdisk25 fcs1 5a primary up secondary hdisk20 fcs0 9b failover up primary hdisk5 fcs0 9a primary </pre>

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 either the Attach Kit `sanlum lun show -p` command or SANpath's `setsp -a` command to determine which host disk device SANpath uses as the I/O access device for the LUN.

Note

For each LUN, SANpath selects one of the path devices and presents it as the single I/O access device for the LUN. This access device does not change, which means it is not necessarily the primary path. You should use this device as a physical volume in LVM. LVM uses it to perform operations on the LUN, and SANpath inhibits access to all the other path devices for this LUN.

To determine this access device for the LUN, look for the line that shows the SANpath device mapping to the device in the output from either the Attach Kit `sanlun lun show -p` command or SANpath's `setsp -a` command. The line should have the following format:

```
spdX = hdiskN
```

X is the number for the SANpath device.

N is the number for the access device. Use this device when you create the volume group.

Using sanlun to find the access device: The following example shows part of the output of the `sanlun lun show -p` command. The SANpath device (spd) for LUN 1 is 3 (hdisk8), and LUN 2 is 4 (hdisk4). This information displays in the first line of output (bold).

SANpath assigns a different SANpath access device for each LUN.

```
sanlun lun show -p
```

```
target1:/vol/vol1/db-lun1 (LUN 1) spd3 = hdisk8
```

```
1g (1073741824) lun state: GOOD
```

```
Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C
```

```
Multipath-provider: Dot Hill SANpath version 4.0.5n
```

host	filer		primary	partner	SANpath	
path	path	device	host	filer	filer	
state	type	filename	HBA	port	port	
					failover	
					priority	
up	secondary	hdisk38	fcs1		5b	failover
up	primary	hdisk33	fcs1	5a		primary
up	primary	hdisk13	fcs0	9a		primary
up	secondary	hdisk8	fcs0		9b	failover

```
target2:/vol/vol1/db-lun2 (LUN 2) spd4 = hdisk4
```

```
100m (104857600) lun state: GOOD
```

```
Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C
```

```
Multipath-provider: Dot Hill SANpath version 4.0.5n
```

host	filer		primary	partner	SANpath	
path	path	device	host	filer	filer	
state	type	filename	HBA	port	port	
					failover	
					priority	
up	secondary	hdisk29	fcs1		5b	failover
up	primary	hdisk24	fcs1	5a		primary
up	secondary	hdisk19	fcs0		9b	failover
up	primary	hdisk4	fcs0	9a		primary

Using setsp -a to find the access device: The following output shows the results of the `setsp -a` command on the same LUNs as the previous example. The SANpath device points to `hdisk8` (`sp3 = hdisk8`) and `hdisk4` (`spd4 = hdisk4`) as the access device for the LUN.

```
# setsp -a
```

```
spd  Path/disk          Status  Pri Exc Buf Balance RtrCnt  RtrDly FailBack
Ockti3xI9Cvf" =====
  3   hdisk13(39,8)      P Good          32    1      9    1000    1
      hdisk33(39,31)     P Good      X
      hdisk8(39,0)       S Good
      hdisk38(39,36)     S Good
spd3 = hdisk8                                ID = "NETAPP LUN"
OckuFYxI19l5" =====
  4   hdisk4(39,7)       P Good      X      32    1      9    1000    1
      hdisk24(39,22)     P Good
      hdisk19(39,17)     S Good
      hdisk29(39,27)     S Good
spd4 = hdisk4                                ID = "NETAPP LUN"
```

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 64, to determine the device to which the SANpath device is mapping.
2	Use the AIX System Management Interface Tool (SMIT) to create a volume group by entering the following command: smit vg

Step	Action																																										
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></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 =Help</td><td></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 colspan="3">F8=Image</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																																											
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><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>																																											
5	Press F10 to exit SMIT.																																										

Step	Action
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 57. ----- 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 NetApp 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.

Step	Action
5	Verify the changes. <pre>/usr/sbin/vxdmpadm stat restored</pre> Result: The command output shows the status of the vxrestore daemon, as shown by the output example below. <pre>The number of daemons running : 1 The interval of daemon: 60 The policy of daemon: check_disabled</pre>

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.

Step	Action																																																								
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. lsdev -Cc disk <table><tr><td>hdisk0</td><td>Available</td><td>1S-08-00-8,0</td><td>16 Bit LVD SCSI Disk Drive</td></tr><tr><td>hdisk1</td><td>Available</td><td>1S-08-00-9,0</td><td>16 Bit LVD SCSI Disk Drive</td></tr><tr><td>hdisk2</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk3</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk4</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk5</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk6</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk7</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk8</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk9</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk10</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk11</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk12</td><td>Available</td><td>1Z-08-01</td><td>NetApp LUN</td></tr><tr><td>hdisk13</td><td>Available</td><td>1D-08-01</td><td>NetApp LUN</td></tr></table>	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	hdisk2	Available	1Z-08-01	NetApp LUN	hdisk3	Available	1Z-08-01	NetApp LUN	hdisk4	Available	1Z-08-01	NetApp LUN	hdisk5	Available	1Z-08-01	NetApp LUN	hdisk6	Available	1D-08-01	NetApp LUN	hdisk7	Available	1D-08-01	NetApp LUN	hdisk8	Available	1D-08-01	NetApp LUN	hdisk9	Available	1D-08-01	NetApp LUN	hdisk10	Available	1Z-08-01	NetApp LUN	hdisk11	Available	1D-08-01	NetApp LUN	hdisk12	Available	1Z-08-01	NetApp LUN	hdisk13	Available	1D-08-01	NetApp LUN
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																																																						
hdisk2	Available	1Z-08-01	NetApp LUN																																																						
hdisk3	Available	1Z-08-01	NetApp LUN																																																						
hdisk4	Available	1Z-08-01	NetApp LUN																																																						
hdisk5	Available	1Z-08-01	NetApp LUN																																																						
hdisk6	Available	1D-08-01	NetApp LUN																																																						
hdisk7	Available	1D-08-01	NetApp LUN																																																						
hdisk8	Available	1D-08-01	NetApp LUN																																																						
hdisk9	Available	1D-08-01	NetApp LUN																																																						
hdisk10	Available	1Z-08-01	NetApp LUN																																																						
hdisk11	Available	1D-08-01	NetApp LUN																																																						
hdisk12	Available	1Z-08-01	NetApp LUN																																																						
hdisk13	Available	1D-08-01	NetApp LUN																																																						
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 NetApp 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 ENCLR_NAME and ENCLR_TYPE. Verify that the ARRAY_TYPE is A/A NETAPP. <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 vxdisksetup. ◆ To add one or more disks to a disk group, enter vx dg add disk 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 filesystem, enter <pre>mount -V vxfs /dev/vx/dsk/diskgroup/volume</pre> To mount the file system, enter <pre>mount -V vxfs /dev/vx/dsk/diskgroup/volume mount_point</pre> smit: Enter smit crvxfs1vstd

Displaying LUN and HBA information

Utility for displaying information

The `sanlun` utility is an attach kit 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 81
- ◆ [“Displaying host HBA information”](#) on page 85

Displaying LUN and HBA information

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 <filer name></code> or the <code>sanlun lun show <filer name:filer 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 1: This abbreviated example displays the type of output you might see if you executed the <code>sanlun lun show</code> command with the <code>-p</code> option on a system with SANpath multipathing. Keep in mind as you look at the output that the following is true: ◆ A path type of primary indicates a direct path to the storage. ◆ A path type of secondary indicates a proxy path through a second storage system. ◆ The target port column always refers to the port on the storage system on which the LUN resides. It always represents a primary path. ◆ Partner port column always refers to the port on the partner storage system. It always represents a secondary path. <pre>sanlun lun show -p</pre> <pre> strife:/vol/vol1/ndu-lun4 (LUN 4) spd8 = hdisk6 100m (104857600) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: A/PG-C Multipath-provider: Dot Hill SANpath version 4.0.5m ----- host filer primary partner SANpath path path device host filer filer failover state type filename HBA port port priority ----- up secondary hdisk31 fcs1 5b failover up primary hdisk26 fcs1 5a primary up secondary hdisk21 fcs0 9b failover up primary hdisk6 fcs0 9a primary </pre>

Step	Action																																																																																																																																																
	Example 2: This example displays sanlun lun show -p output for a system with VERITAS multipathing. sanlun lun show -p 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 <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>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> 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 <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>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 LUN and HBA information

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 of what you need to do to configure SAN booting on an AIX host, and also provides detailed procedures.

Note

SAN booting is currently not supported with VERITAS configurations. Network Appliance qualifies solutions and components. To verify that SAN booting is supported in your configuration, see the *FCP Host Compatibility Matrix* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/fcp_support.shtml.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[SAN boot overview](#)” on page 88
- ◆ “[Creating the boot LUN](#)” on page 90
- ◆ “[Booting the AIX Installation CD-ROM](#)” on page 94
- ◆ “[Installing the OS](#)” on page 95
- ◆ “[Troubleshooting the SAN boot LUN](#)” on page 102

SAN boot overview

What is SAN booting

SAN booting is using a SAN-attached disk, such as a LUN, as a boot device for a host. Configuring SAN booting allows you to do the following:

- ◆ Remove the hard drives from your servers and use the SAN for your boot needs, thereby eliminating the costs associated with maintaining and servicing hard drives.
- ◆ Use the SAN for both storage and boot needs, thereby consolidating and centralizing storage.

Methods for creating a bootable LUN

This chapter provides a procedural example that uses the AIX Base Operating System to configure a SAN boot LUN and install a new OS image onto it. The procedure is intended as a detailed example for creating a bootable LUN that can be applied to most configurations. Other methods for creating a bootable LUN, such as using the AIX Alternate Disk Install, which allows you to install the OS while it is up and running, are not described. Refer to AIX documentation for details about other installation methods.

Pre-requisites

The examples in this chapter assume the following:

- ◆ The host is booting from an existing internal OS disk. It is not a diskless server.
- ◆ The FCP IBM AIX Attach Kit 1.3 is installed on the host.

Configuration overview

Configuring SAN booting involves:

1. Configuring the igroup and SAN boot LUN.
 - a. Determining the size of the boot LUN.
 - a. Creating the igroup.
 - b. Creating the boot LUN.
 - c. Mapping the boot LUN to the igroup.
2. Installing the OS on the boot LUN.

- a. Restarting the system and booting the AIX installation CD-ROM.
- b. Selecting the type of install.
- c. Determining the primary boot path
- d. Installing the OS onto the boot LUN.

**For detailed
information**

The following sections provide details about creating a bootable SAN LUN:

- ◆ “[Determining the size of the boot LUN](#)” on page 90
- ◆ “[Configuring the igroup and boot LUN](#)” on page 90
- ◆ “[Restarting the host](#)” on page 94
- ◆ “[Selecting the type of Installation](#)” on page 95
- ◆ “[Determining the primary boot path and installing the OS](#)” on page 98

Configuring the igroup and LUN

Creating the boot LUN

Determining the size of the boot LUN

You must create a LUN that is large enough to provide space for the AIX version you are using. AIX product documentation specifies a minimum of 2.2 GB of disk space are required for AIX 5.1, 5.2, and 5.3 operating systems.

Configuring the igroup and boot LUN

After you determine the size of the boot LUN, you must create the igroup, create the bootable LUN, and map the bootable LUN to the igroup.

Creating the igroup: To create the igroup, complete the following steps. If you already have an existing AIX igroup that contains the WWPNs of the host’s HBAs, proceed to “[Creating and mapping the boot LUN](#)” on page 91.

Step	Action
1	Log in as root on the host.
2	Get the WWPN for each of the host’s HBAs by entering the following <code>sanlun</code> command. The <code>sanlun</code> command is in the <code>/opt/NetApp/santools/bin</code> directory. sanlun fcp show adapter -c Result: The <code>-c</code> option tells the <code>sanlun</code> command to generate on the host console a complete <code>igroup create</code> command that includes the host’s WWPNs. The igroup’s name in the <code>igroup create</code> command matches the name of the host. You can copy this command from the host console and use it to create the igroup in Step 4. Example: The <code>sanlun fcp show adapter -c</code> command produces the following output on the console of an AIX system called <code>boothost</code>. This example uses two WWPNs. Enter this filer command to create an initiator group for this system: <pre>igroup create -f -t aix "boothost" 100a0b000010a700 100a0b000010a606</pre> Because the name of the host is “boothost,” this command line produces an igroup called “boothost.”

Step	Action
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 "boothost" 100a0b000010a700 100a0b000010a601</pre>
5	Paste the command on the storage system command line and press Enter. Result: The storage system creates the <code>igroup</code> and gives it the same name as the host name.
6	On the storage system command line, verify that the new <code>igroup</code> is created by entering the following command: igroup show Result: The console displays the newly created <code>igroup</code> with the host's WWPNs. The following is output from the <code>igroup show</code> command displays. <pre>storage_1> igroup show -v boothost boothost (FCP): OS Type: aix Member: 10:0a:0b:00:00:10:a7:00 Member: 10:0a:0b:00:00:10:a6:01</pre>

Creating and mapping the boot LUN: To create the boot LUN and map it to an `igroup`, complete the following steps:

Step	Action
1	On the storage system command line, enter the following command: <pre>lun create -s size -t ostype lun-path</pre> size is the size of the LUN, which you obtained in “Determining the size of the boot LUN” on page 90. You can use one of the following characters to establish the size: c (bytes), w (words or double bytes), b (512), k (kilobytes), m (megabytes), g (gigabytes), t (terabytes). -t ostype indicates the type of operating system running on the host that accesses the boot LUN. Enter aix. <i>lun-path</i> is the LUN’s path name that includes the volume and qtree. Example: <pre>lun create -s 18g -t aix /vol/bootvol/bootlun1</pre>
2	Enter the following command to map the LUN to the initiator group. <pre>lun map lun-path igroup-name LUN-ID</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. <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. The default LUN ID begins with 0 and increments by 1 for each additional LUN that is created. If you have existing LUNs, Data ONTAP assigns the next available LUN ID. Example: <pre>lun map /vol/bootvol/bootlun1 boothost 0</pre>

Step	Action
3	Verify that the LUN is online and mapped to the igroup by entering the following command: <pre>lun show -v</pre> Example: <pre>storage_1> lun show -v /vol/bootvol/bootlun1 18g (19327352832) (r/w,online, mapped) Serial#: OdCtSYv7Y6w3 Share: none Space Reservation: enabled Multiprotocol Type: aix Maps: boothost=0</pre>

Installing the OS on the boot LUN

Booting the AIX Installation CD-ROM

Restarting the host You must restart the host using the AIX CD-ROM to activate the AIX Base Operating System Install. The AIX Base Operating System Install contains the installation and maintenance scripts that enable you to install the AIX OS onto the boot LUN.

Complete the following steps to restart the host and activate the AIX Base Operating System Install.

Step	Action
1	Insert the AIX CD-ROM into the drive of the host system.
2	Enter the following to change the bootlist on the host system to the CD-ROM drive: <code>bootlist -m normal <cd_drive></code> Example: The following command changes the bootlist to the first CD-ROM drive. <code>bootlist -m normal cd0</code>
3	Enter <code>shutdown -Fr now</code> to shut down and restart the host system. Result: The system reboots from the AIX CD-ROM.
4	At the Type a 1 and press Enter to use this terminal as a system console, enter 1 .
5	At the Type 1 and press enter to have English during the install, enter 1 to proceed in English. Result: The system displays the Base Operating System Installation and Maintenance screen.

Installing the OS on the boot LUN

Installing the OS

Selecting the type of Installation

Use the AIX Base Operating System, and follow the steps below to select the type of OS you want to copy.

Step	Action
1	At the Installation and Maintenance screen, enter 2 to select Change/Show Installation Settings and Install Result: The system displays Installation and Settings options.
2	Enter 2 to select Change/Show Installation Settings and Install. Result: The system displays additional Installation and Settings options.
3	Enter 1 to select System Settings. Result: The system displays OS installation options.
4	Take one of the following options: ◆ Enter 1 to select New and Complete Overwrite. ◆ Enter 2 to select Preservation Install. ◆ Enter 3 to select Migration Install.
5	Enter 77 to navigate through Display More Disk information screens until the system displays the Device Adapter Connection information shown in the example that follows. Keep this display available, and go to “ Determining the primary boot path and installing the OS ” on page 98.

Example

The following example selects a New and Complete Overwrite, and then enters 77 to display Device Adapter Connection Information.

```
Welcome to Base Operating System
Installation and Maintenance
```

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

```
>>> 1 Start Install Now with Default Settings
      2 Change/Show Installation Settings and Install
      3 Start Maintenance Mode for System Recovery

      88 Help ?
      99 Previous Menu
```

```
>>> Choice [1]: 2
```

Installation and Settings

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

```
1 System Settings:
    Method of Installation.....Preservation
    Disk Where You Want to Install....hdisk0

2 Primary Language Environment Settings (AFTER Install):
    Cultural Convention.....English (United States)
    Language .....English (United States)
    Keyboard .....English(United States)
    Keyboard Type.....Default

3 More Options (Desktop, Security, Kernel, Software, ...)
```

```
>>> 0 Install with the current settings listed above.
```

```
+-----
WARNING: Base Operating System Installation will destroy or impair
recovery of SOME data on the destination disk hdisk0.
-----
```

```
      88 Help ?
      99 Previous Menu
>>> Choice [0]: 1
```

Change Method of Installation

Type the number of the installation method and press Enter.

```
1 New and Complete Overwrite
    Overwrites EVERYTHING on the disk selected for installation.
    Warning: Only use this method if the disk is totally empty or if
    there is nothing on the disk you want to preserve.
```

```
>>> 2 Preservation Install
Preserves SOME of the existing data on the disk selected for
installation. Warning: This method overwrites the usr (/usr),
variable (/var), temporary (/tmp), and root (/) file systems.
Other product (applications) files and configuration data will
be destroyed.

3 Migration Install
Upgrades the Base Operating System to the current release.
Other product (applications) files and configuration data are
saved.

88 Help ?
99 Previous Menu
```

```
>>> Choice [2]: 1
```

Change Disk(s) Where You Want to Install

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

	Name	Location Code	Size(MB)	VG Status	Bootable
>>> 1	hdisk0	1S-08-00-5,0	34715	rootvg	Yes No
2	hdisk1	1S-08-00-8,0	34715	rootvg	Yes No
3	hdisk2	1Z-08-01	36864	none	Yes No
4	hdisk3	1Z-08-01	12288	none	Yes No
5	hdisk5	1Z-08-01	12288	none	Yes No
06	MORE CHOICES...				
>>> 0	Continue with choices indicated above				
66	Devices not known to Base Operating System Installation				
77	Display More Disk Information				
88	Help ?				
99	Previous Menu				

```
>>> Choice [0]: 77
```

Change Disk(s) Where You Want to Install

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

Name	Physical Volume Identifier
------	----------------------------

```
>>> 1  hdisk0    0000dffa5334e3df
      2  hdisk1    0000dffa6felaf37
      3  hdisk2    0000dffa1c7c3ab9
      4  hdisk3    0000000000000000
      5  hdisk5    0000000000000000
      06  MORE CHOICES...
```

```
>>> 0  Continue with choices indicated above
```

```
66  Devices not known to Base Operating System Installation
77  Display More Disk Information
88  Help ?
99  Previous Menu
```

```
>>> Choice [0]: 77
```

Change Disk(s) Where You Want to Install

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

Name	Device Adapter Connection Location or Physical Location Code
------	---

```
>>> 1  hdisk0    none
      2  hdisk1    none
      3  hdisk2    0x50a9800302008da0//0x0
      4  hdisk3    0x50a9800302008da0//0x10000000000000
      5  hdisk5    0x50a9800c02008da0//0x10000000000000
      06  MORE CHOICES...
```

Determining the primary boot path and installing the OS

During this procedure the Base Operating System prompts you for an adapter connection to the boot LUN. The connection that you specify must be a primary path to the LUN.

Follow the steps below to determine the primary path and configure the adapter connection location.

Step	Action														
1	Keep the Device Adapter Connection Location display available. Example: The Device Adapter Connection Location screen displays hdisk connection information similar to the following: <pre>Change Disk(s) Where You Want to Install</pre> Type one or more numbers for the disk(s) to be used for installation and press Enter. To cancel a choice, type the corresponding number and Press Enter. At least one bootable disk must be selected. The current choice is indicated by >>>. <table><thead><tr><th>Name</th><th>Device Adapter Connection Location or Physical Location Code</th></tr></thead><tbody><tr><td>>>> 1 hdisk0</td><td>none</td></tr><tr><td>2 hdisk1</td><td>none</td></tr><tr><td>3 hdisk2</td><td>0x50a9800302008da0//0x0</td></tr><tr><td>4 hdisk3</td><td>0x50a9800302008da0//0x1000000000000</td></tr><tr><td>5 hdisk5</td><td>0x50a9800c02008da0//0x1000000000000</td></tr><tr><td>06</td><td>MORE CHOICES...</td></tr></tbody></table>	Name	Device Adapter Connection Location or Physical Location Code	>>> 1 hdisk0	none	2 hdisk1	none	3 hdisk2	0x50a9800302008da0//0x0	4 hdisk3	0x50a9800302008da0//0x1000000000000	5 hdisk5	0x50a9800c02008da0//0x1000000000000	06	MORE CHOICES...
Name	Device Adapter Connection Location or Physical Location Code														
>>> 1 hdisk0	none														
2 hdisk1	none														
3 hdisk2	0x50a9800302008da0//0x0														
4 hdisk3	0x50a9800302008da0//0x1000000000000														
5 hdisk5	0x50a9800c02008da0//0x1000000000000														
06	MORE CHOICES...														
2	Use a telnet command to access the storage system.														

Step	Action
3	Enter fcv show adapter and locate the FCP portname that matches the portname displayed in the Device Adapter Connection Location screen. Example: The system displays information similar to the following. In the following example, the portname for Slot 5a (local) matches the portname for hdisk3 on the Device Adapter Connection Location screen, shown in previous step. This is the primary path. <pre> fcv show adapter Slot: 5a Description: Fibre Channel Target Adapter 5a Dual-channel, QLogic 2312 (2342) rev. 2) Adapter Type: Local Status: ONLINE FC Nodename: 50:a9:80:00:02:00:8d:a0 (50a9800002008da0) FC Portname: 50:a9:80:03:02:00:8d:a0 (50a9800302008da0) Standby: No Slot: 5b Description: Fibre Channel Target Adapter 5b (Dual-channel, QLogic 2312 (2342) rev. 2) Adapter Type: Partner Status: ONLINE FC Nodename: 50:a9:80:00:02:00:8d:b8 (50a9800002008db8) FC Portname: 50:a9:80:0c:02:00:8d:b8 (50a9800c02008db8) Standby: No </pre>
4	On the Device Adapter Connection Location screen, enter the number for the disk and primary path that you identified in the previous step. Result: The system indicates the disk you selected and asks if you want to continue.
5	Enter 0 at the Choice prompt to confirm your selection. Result: The system displays the Installations and Setting screen and asks you to confirm your installation choices.
6	Enter 0 at the Choice prompt to install the OS onto the boot LUN. Result: The system displays Overwrite Installation Summary information and asks you to confirm your installation choices.

Step	Action
7	Enter 1 at the <code>Choice</code> prompt to continue with the install. Result: The system copies the OS from the CD-ROM onto the boot LUN and uses the SAN boot LUN to reboot the system.

Troubleshooting the SAN boot LUN

If you cannot boot from the primary or alternate path

If the system does not boot from the primary or alternate path, check the following configuration components:

- ◆ Connectivity—Check physical connectivity and cabling between the HBA, the switches, and the cluster.
- ◆ Boot path settings—If your system fails to boot from either the primary or alternate path, boot from the internal system disk, and then check your boot path settings.

Checking boot path settings: To boot from the system disk and check your boot path settings, complete the following steps:

Step	Action
1	If necessary, reset the system to return to the SMS prompt.
2	Boot from the internal system disk using the SMS menu to select an internal system boot disk.
3	Verify that you correctly translated the primary and secondary device paths to the boot LUN to the correct hardware paths. Enter sanlun lun show -p Result: ◆The sanlun lun -p command displays the primary and the primary and secondary device paths to the boot LUN and also shows the hardware paths to the device files.
4	Make sure that the path to the boot LUN is in the root volume group by entering the following command: lsvg rootvg Result: Paths to the LUN are listed under the <code>Physical volumes</code> section of the output.

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 104
- ◆ [“Installation fails because system does not meet requirements”](#) on page 106
- ◆ [“Commands you can use to troubleshoot problems”](#) on page 110
- ◆ [“Use SANpath commands when working with multipathing”](#) on page 112
- ◆ [“Attach Kit diagnostic utilities”](#) on page 114
- ◆ [“Change the AIX HBA’s attributes”](#) on page 119
- ◆ [“Change the HBA queue depth with no LUNs mapped”](#) on page 124
- ◆ [“Change the value for the queue depths on the disk”](#) on page 125
- ◆ [“General troubleshooting”](#) on page 126

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 filer FCP cfmodes is correctly set for your configuration. See Chapter 4, “[Migrating a Storage Cluster to Single Image Mode](#),” on page 33.
- ◆ 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 `setsp -a` 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 Attach Kit 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 SANpath or 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

Install command checks for system requirements

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

Rerunning the installation script

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	Rerun the install script. Result: The script normally resumes at the point where it failed and runs until completion. Note _____ Using the uninstall command to remove the FCP IBM AIX Attach Kit software is not necessary. _____

Example of a failed installation

This example shows an installation that fails because a package was not present. For this example, the following is true:

- ◆ The Attach Kit software was downloaded to /tmp.
- ◆ The SANpath package was downloaded and copied to the /tmp/SANpath directory.
- ◆ The install script is extracting files to /tmp/ntap_aix_fcp_1.3.
- ◆ The installation fails because the HBA API is missing.

Note

This output has been shortened.

```
# cd /tmp/ntap_aix_fcp_1.3
# ./install
Checking HBAs on the system... done.

Please enter the location of the SANpath package:
(if you are installing from CD, you may enter the device name, e.g.
/dev/cd0)
[default: /tmp/ntap_aix_fcp_1.3]

> /tmp/SANpath

Installing SANpath from /tmp/SANpath... done.

In order to enable SANpath, you must enter your authorization data.
Enter your company name: > MyCompany
Enter SANpath serial number: > XXXX
Enter authorization code: > XXXX

SANpath authorization accepted.

Please enter the location of the NetApp packages:
(if you are installing from CD, you may enter the device name, e.g.
/dev/cd0)
[default: /tmp/ntap_aix_fcp_1.3]
>

Installing NetApp packages from /tmp/ntap_aix_fcp_1.3... failed!
ERROR: NetApp packages failed to install due to missing
prerequisites.
Install the prerequisites and run this script again to resume
installation.
See your documentation and the installation log for details.

Installation log saved in /tmp/netapp_install_log_031208151114

To confirm which prerequisite is missing, check the log file
/tmp/netapp_install_log_031208151114. The following is the section of the log
file that tells you which prerequisite is missing.

...
Verifying requisites...done
Results...
```

FAILURES

Filesets listed in this section failed pre-installation verification and will not be installed.

Requisite Failures

SELECTED FILESETS: The following is a list of filesets that you asked to install. They cannot be installed until all of their requisite filesets are also installed. See subsequent lists for details of requisites.

```
#1 NetApp.fcp_toolkit.sanlun 1.3.0.0 # Network Appliance FCP
Toolki...
```

MISSING REQUISITES: The following filesets are required by one or more of the selected filesets listed above. They are not currently installed and could not be found on the installation media. (Selected filesets which depend upon these requisites are referenced in parentheses.)

```
devices.common.IBM.fc.hba-api 5.1.0.0      # Base Level Fileset
      (dep #s: 1)
```

...

Example of restarting the installation

Once you install the missing fileset, restart the installation. It resumes where it left off. If there are no other problems, it completes successfully.

Note

Running the uninstall script before you complete the installation is not necessary.

The following is an example of the output you would see if you corrected the problem in the previous example and then restarted the installation.

Note

This output has been shortened.

```
# ./install
Checking HBAs on the system... done.
SANpath is already installed.
```

Please enter the location of the NetApp packages:

```
(if you are installing from CD, you may enter the device name, e.g.  
/dev/cd0)  
[default: /tmp/ntap_aix_fcp_1.3]
```

```
>
```

```
Installing NetApp packages from /tmp/ntap_aix_fcp_1.3... done.  
HBA settings update successful.  
Running bosboot to update the boot image... done.  
Installation of the NetApp AIX Attach Kit was successful.  
You MUST reboot the system before accessing your NetApp filer(s).  
  
Installation log saved in both /tmp/netapp_install_log_031208161356  
and /opt/NetApp/install_log
```

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 Attach Kit `sanlun` utility
- ◆ The AIX `lsdev` command
- ◆ SANpath commands `sppath -a` and `setsp -a` (see [“Use SANpath commands when working with multipathing”](#) on page 112)
- ◆ The VERITAS commands `vxddm adm` 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 56.

Note

You can use both the Attach Kit `sanlun`, Dot Hill’s SANpath or DMP to check your configuration. By comparing the output from `sanlun`, SANpath or DMP, you can verify that both of these utilities are working correctly.

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 -C -c disk
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
hdisk2 Available 1n-08-01 NetApp LUN
hdisk3 Available 1n-08-01 NetApp LUN
hdisk4 Available 1n-08-01 NetApp LUN
hdisk5 Available 1n-08-01 NetApp LUN
hdisk6 Available 1n-08-01 NetApp LUN
hdisk7 Available 1n-08-01 NetApp LUN
hdisk8 Available 1n-08-01 NetApp LUN
hdisk9 Available 1n-08-01 NetApp LUN
hdisk10 Available 1n-08-01 NetApp LUN
hdisk11 Available 1n-08-01 NetApp LUN
hdisk12 Available 1n-08-01 NetApp LUN
hdisk13 Available 1n-08-01 NetApp LUN
hdisk14 Available 1n-08-01 NetApp LUN
hdisk15 Available 1n-08-01 NetApp LUN
hdisk16 Available 1n-08-01 NetApp LUN
hdisk17 Available 1n-08-01 NetApp LUN
hdisk18 Available 1n-08-01 NetApp LUN
hdisk19 Available 1n-08-01 NetApp LUN
hdisk20 Available 1n-08-01 NetApp LUN
hdisk21 Available 1n-08-01 NetApp LUN
...
hdisk80 Available 11-08-01 NetApp LUN
hdisk81 Available 11-08-01 NetApp LUN
hdisk82 Available 11-08-01 NetApp LUN
hdisk83 Available 11-08-01 NetApp LUN
hdisk84 Available 11-08-01 NetApp LUN
hdisk85 Available 11-08-01 NetApp LUN
hdisk86 Available 11-08-01 NetApp LUN
hdisk87 Available 11-08-01 NetApp LUN
hdisk88 Available 11-08-01 NetApp LUN
hdisk89 Available 11-08-01 NetApp LUN
hdisk90 Available 11-08-01 NetApp LUN
hdisk91 Available 11-08-01 NetApp LUN
hdisk92 Available 11-08-01 NetApp LUN
hdisk93 Available 11-08-01 NetApp LUN
hdisk94 Available 11-08-01 NetApp LUN
hdisk95 Available 11-08-01 NetApp LUN
hdisk96 Available 11-08-01 NetApp LUN
hdisk97 Available 11-08-01 NetApp LUN
```

Use SANpath commands when working with multipathing

Use SANpath to check paths for multipathing

If you have a SANpath configuration, you can use SANpath to display information about the paths to the devices and data.

This information is important because the order of the paths you use when you set up the LUNs can affect performance. In general, secondary paths are slower than primary paths. So, to get the best performance, you must set up your paths with the primary paths first and then follow them with the secondary paths.

Use the SANpath `sppath -a` and `setsp -a` commands to see the path ordering for your system and correct it, if necessary.

For detailed information on what you can do with SANpath and how to use its commands, see the *SANpath User's Guide for IBM AIX Operating Systems* that comes with the multipathing version of the Attach Kit. A description of this document is in the section [“Where to go for more information”](#) on page 11.

Path information available from the `setsp -a` command

SANpath's `setsp -a` command displays columns with the following information:

- ◆ `spd` is the LUN's SANpath device number.
- ◆ `Path/disk` lists the device names and disk numbers for each of the redundant physical paths to the device. The disk number is appended to “`hdisk.`”
- ◆ `Status` tells you whether the path is primary (P) or secondary (S) and whether it is functioning normally (Good) or has failed (Bad)
- ◆ `Pri` is the Primary field. It indicates the one primary path that is the highest primary path used when load balancing is disabled. There is always only one path selected in this field.
- ◆ `Buf` is the number of buffer pointers (or buffer structures) pre-allocated for each logical device.
- ◆ `Balance` indicates whether dynamic load balancing is enabled for the device (0 indicates disabled; 1 indicates enabled).
- ◆ `RtrCnt` is the number of times SANpath retries the I/O operation on an arbitrary secondary path after an I/O operation has failed on all of the paths. When the specified number of retries has failed, the I/O fails.
- ◆ `RtrDly` is the time interval, in milliseconds, between the retry attempts described in the `RtrCnt`.

- ◆ Failback indicates whether failback is enabled for the logical device. If so, SANpath keeps testing a path that has failed until it can be returned to service.

Example of SANpath's setsp -a output

When you enter SANpath's `setsp -a` command, it displays the following information:

```
# setsp -a
```

spd	Path/disk	Status	Pri	Exc	Buf	Balance	RtrCnt	RtrDly	FailBack
1	hdisk6(38,4)	P Good			32	1	9	1000	1
	hdisk8(38,5)	P Good	X						
	hdisk7(38,6)	S Good							
	hdisk9(38,7)	S Good							
spd1 = hdisk6		ID = "NETAPP LUN Vqi5S3mKMsyD"							

setsp -v provides information on system limits

SANpath permits up to 128 LUNs and a maximum of 32 paths to each LUN. If your system configuration exceeds these limits, the following warning appears at the end of the `setsp -v` output.

```
setsp: WARNING: SANpath detected more than 128 SANpath-compatible
devices (LUNs), or more than 32 paths to a LUN. SANpath cannot
protect more than 128 LUNs or more than 32 paths to a LUN. Please
modify your configuration so that no more than 128 LUNs are
available to this host, and no more than 32 paths are available to
any given LUN. "setsp -g" will show the detailed configuration
error.
```

If you receive this warning, you must correct your system configuration to ensure that it is within the SANpath limits.

Attach Kit diagnostic utilities

Diagnostic utilities The Attach Kit 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 115 for more information.
- ◆ `aix_info`. See “[Using the AIX data-collection tool](#)” on page 116.
- ◆ `<switch>_info`, where `<switch>` is the brand of switch you are using (for example, `brocade_info`). See “[Collecting switch information](#)” on page 117.

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 <code>/opt/NetApp/santools/bin</code> : <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> : <code>cd /opt/NetApp/santools/bin</code>
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: <pre>cd /opt/NetApp/santools/bin</pre>
2	At the command prompt, enter the following command: <pre>aix_info [-d <directory>] [-n <file_name>]</pre> <code>-d <directory></code> specifies the location for the gzipped file that <code>aix_info</code> creates. The default location is <code>/tmp/netapp</code>. <code>-n <file_name></code> specifies the name for the output file and directory names. The default pathname is <code>/tmp/netapp/ntap_aix_info.tar.gz</code>. Note _____ Ensure that you have at least 64MB of free storage in the output directory Result: When you run <code>aix_info</code>, 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 <code>errpt</code> command. Example: Depending on the options you specify, you see output similar to the following when you run this utility: <pre># aix_info AIX system info is in directory /tmp/netapp/ntap_aix_info Compressed file is /tmp/netapp/ntap_aix_info.tar.gz 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 Attach Kit supplies when you install the FCP IBM AIX Attach Kit software. The Attach Kit provides a utility for each switch it supports for this Attach Kit. 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 path_and_file_name 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 user:password 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 switchname, where switchname 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 switch_brand is the brand of the switch, such as Brocade, and switchname 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 124).

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	If you have SANpath installed, you must get the SPD numbers for all SANpath devices where at least one path is a child device of the FCP device. Use either SANpath's <code>setsp -v</code> command or the Attach Kit <code>sanlun lun show -p</code> to get this information.

Step	Action
4	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.
5	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.
6	Disable any applications that are accessing logical volumes in the volume groups.
7	Vary off all the volume groups and unmount any file systems. Example: To vary off the volume group called <code>my_vg</code>, you would enter the following command. <pre># varyoffvg my_vg</pre>
8	If you have SANpath installed, stop the SPD devices you found in step 3. Example: This example halts the <code>sp0</code> device. <pre># setsp -T -l0 spd0: stop device OK.</pre>
9	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>

Step	Action
10	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>
11	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>
12	Reload the HBA device to enable the new setting. Example: This example reloads the FCP device fscsi1. <pre># mkdev -l fcs1 fcs1 Available</pre>
13	If necessary, reload the FCP device fscsi1. Example: The following command line reloads the FCP device fscsi1. <pre># mkdev -l fscsi1 fscsi1 Available</pre>
14	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>

Step	Action
15	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>
16	Remount any file systems that you unmounted.
17	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 120.

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 install script 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

- access device 65
- Array Support Library
 - see ASL 28

ASL

- definition 28

attributes

- changing and rebooting 119
- changing on HBA 119
- changing with SMIT 119
- changing without rebooting 120

C

CD-ROM

- getting software 16

- cfgmgr command 56, 63, 74

cfmode

- correct setting 33

- chdev command 119

checklists

- filer setup 104
- host setup 104
- installation tasks 2
- troubleshooting 104

commands

- igroup create 91
- igroup show 91

configuration

- direct-attached 48

D

diagnostic utilities

- aix_info 114
- executing 114
- filer_info 114, 115
- switch_info 114, 117

direct-attached configuration

- set media type 48

disk

- changing value of queue depths 125

disk shelves

- powering up 10

E

equipment

- powering up 10

F

file system

- creating 59, 69

filer

- correct dfmode 33
- defined v
- direct-attached 48
- powering up 10
- setup checklist 104
- verifying setup 8

filer_info

- executing 115

files

- filer_info 115
- switch_info 117

fileset

- check version 7

H

HBA

- changing attributes 119
- changing queue depth 124
- installing 8
- verifying queue depth 21

host 14

- check OS version 7
- creating volume group 57, 67
- device for physical volume 64
- discovering new LUNs 56, 63, 74
- displaying LUN information 81
- powering up 10
- requirements 6
- SANpath access device 65

setup checklist 104

I

igroup

- create command 91

- verify 91

igroup create command 91

igroup show command 51, 91

install command

- checks system requirements 106

- example (restarting) 108

- rerunning 106

- running interactively 19

- what it does 18

installation

- failed example 106

- installing host software 18

- pre-installation tasks 6

- restarting 108

- uninstalling the host software 23

L

lsdev command 120

- using for troubleshooting 111

lslpp command 7

lun commands

- lun create 52, 92

- lun map 52, 92

LUN ID

- defined v

lun setup command 51

lun show 53

LUNs

- best practices for creating 49

- creating with lun setup 51

- defined v

- discovering 56, 63, 74

- displaying host information 81

- tasks for creating on filer 49

- verifying 53

M

media type

- setting 48

multipathing

- SANpath access device 65

N

NOW site

- downloading software 15

O

oslevel command 7

P

physical volume

- selecting device 64

powering up equipment 10

Q

queue depth

- changing value on disk 125

- changing value on HBA 124

- verifying 21

R

requirements 6

S

sanlun fcp show adapter command 50, 85

sanlun lun show command 81

sanlun utility

- about 80

- sanlun fcp show adapter 85

- using for troubleshooting 110

SANpath

- checking paths 112

- choosing access device 65

- setsp -a command 112

- setsp -v command 113

- setsp -a
 - SANpath command 112
- setsp -v
 - SANpath command 113
- SMIT
 - creating file system 59, 69
 - creating volume group 57, 67
 - unloading child devices 119
- software
 - downloading 15
 - getting 14
 - getting from CD-ROM 16
 - installation tasks 18
 - uncompressing 16
- storage
 - accessing on volume group 59, 69
- switch_info
 - executing 117
- switches
 - powering up 10
 - verifying setup 9
 - zoning 9

system requirements 6

T

- Terminology v
- troubleshooting
 - checklists 104
 - using lsdev command 111
 - using sanlun 110

U

- uninstall command
 - example (command-line mode) 25
 - using interactively 23
 - what it does 22

V

- volume group
 - accessing storage 59, 69
 - creating 57, 67

Z

- zoning 9

