

FCP AIX® Host Utilities 4.1 for Native OS 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 215-03246_A0
June 2007

Copyright and trademark information

Copyright information

Copyright © 1994–2007 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL NETWORK APPLIANCE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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

StoreVault, vFiler, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the U.S.A.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or 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	vii
Chapter 1	Host Utilities Installation and Setup Overview	1
	Host Utilities overview	2
	FCP AIX Host Utilities installation overview	7
	Where to go for more information	10
Chapter 2	Installing the Host Utilities Software.	13
	Completing the pre-installation tasks.	14
	Obtaining the software packages	15
	Installing the Host Utilities software packages.	18
	Uninstalling the software packages.	32
Chapter 3	Migrating a Storage Cluster to Single Image Mode.	39
	Deciding whether to migrate to single image mode	40
	Understanding cfmode	41
	Tasks to complete before you change the cfmode	45
	Upgrading to single image mode	47
Chapter 4	Creating, Configuring, and Managing LUNs	53
	Overview of completing the configuration	54
	Creating and mapping igroups and LUNs	55
	LUN configuration overview.	57
	Configuring LUNs with native AIX LVM.	60
	Configuring LUNs with MPIO software on non-VIO hosts	62
	Configuring LUNs with MPIO software on VIO servers	66
	Creating volume groups and file systems on AIX hosts	73
	Using sanlun to display host LUN information	79
	Displaying host HBA information	82

Chapter 5	Configuring SAN booting.	85
	SAN boot overview.	86
	Setting up a SAN Boot LUN	89
	Cloning the root volume group	94
 Chapter 6	 Troubleshooting	 97
	Troubleshooting checklists	98
	Confirming that your system is correctly set up	100
	Installation fails because system does not meet requirements	106
	Commands you can use to troubleshoot problems	107
	SAN Toolkit diagnostic utilities	108
	Change the AIX HBA's attributes	114
	Change the value for the queue depths on the disk.	118
	General troubleshooting	119
	 Index	 121

Preface

About this guide

This guide describes how to integrate an AIX® host with a Network Appliance® storage system with Data ONTAP® software and the Fibre Channel Protocol (FCP) for SCSI.

See either the *NetApp Interoperability Matrix Tool*, available on the NetApp on the Web™ (NOW) site <http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp FCP and iSCSI Products*, available at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>, for current information about the system requirements.

This document applies only to host environments running Multiple Path I/O (MPIO). In addition, it supports host environments running the AIX Virtual I/O Server (VIO). If you are using Veritas® as your multipathing solution, see the *FCP AXI Host Utilities for Veritas Installation and Setup Guide*.

Audience

This guide is for system installers and administrators who:

- ◆ Are familiar with the AIX operating system
- ◆ Are familiar with MPIO
- ◆ Are familiar with VIO server environments, if that is the environment being used
- ◆ Understand how to configure and administer a storage system
- ◆ Understand how to use FCP to store and retrieve data on a storage system.

Terminology

This guide uses the following terms:

- ◆ NetApp® storage products (filers, FAS appliances, and NearStore systems) are all storage systems and are also sometimes called filers or storage appliances.
- ◆ *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.

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.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fctest <i>adapternam</i></code> command, you enter the characters “fctest” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ 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 your program is case-sensitive and uppercase letters are necessary for it to work properly.

Keyboard conventions

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it means...
hyphen (-)	Used to separate individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter, enter</i>	Used to refer to the key that generates a carriage return; the key is named Return on some keyboards. Also used to mean pressing one or more keys and then pressing the Enter key.
<i>type</i>	Used to mean pressing one or more keys on the keyboard.

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.

How to send your comments

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this publication, you can send us comments electronically to doccomments@netapp.com. To help us direct the comments to the correct division, please put FCP AIX in the subject line.

About this chapter

This chapter provides a brief overview of the FCP AIX Host Utilities for Native OS.

In addition, this chapter provides a high-level overview of the tasks you need to perform to install the Host Utilities and set up your system. It also includes pointers to other documentation that might be helpful in working with your Host Utilities setup.

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Host Utilities overview”](#) on page 2
- ◆ [“FCP AIX Host Utilities installation overview”](#) on page 7
- ◆ [“Where to go for more information”](#) on page 10

Host Utilities overview

Important information about Host Utilities and Attach Kits

Previously NetApp provided FCP AIX Attach Kits to enable users to connect AIX hosts to NetApp storage systems. Going forward, NetApp is using Host Utilities to provide this support.

Host Utilities both simplify the products and provide more options for the host stacks. A single FCP AIX Host Utilities can contain software tools and documentation packages to support different host configurations.

Because the Host Utilities are the umbrella for the different kits, component version numbers are not tied to the Host Utilities version number. It is possible for each component to have a different version number. For example, at the time this document was produced, it supported the FCP Host Utilities 4.1 using version 3.1 of the SAN Toolkit.

If you are concerned about the version numbers, you can confirm that you have the most current components by checking the Host Utilities version number in the *NetApp Interoperability Matrix Tool* at

<http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at

<http://now.netapp.com/NOW/knowledge/docs/docs.shtml> *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Overview of the FCP AIX Host Utilities

The Host Utilities provide the tools and information necessary to enable an AIX host to connect to your storage systems. The FCP AIX Host Utilities for Native OS use Multiple Path I/O (MPIO) with IBM's default Path Control Module (PCM) as the multipathing solution. In addition, Host Utilities work with hosts running the AIX Virtual I/O (VIO) server.

The Host Utilities consist of:

- ◆ The SAN Toolkit, which contains:
 - ❖ Configuration tools that help you configure host, HBA, and system files and create persistent bindings.

You can configure the AIX host using the tools provided in the FCP AIX Host Utilities, HBA provided tools, or a combination FCP AIX Host Utilities and HBA provided tools.

- ❖ The `sanlun` utility that helps you manage LUNs and the host HBA.
- ❖ Diagnostic scripts that provide diagnostic information about components in your configuration. Customer Support may ask for output from these scripts.
- ◆ Documentation that describes how to configure the AIX host. It includes this setup guide, *Release Notes*, and a quick command reference guide.

The most current information on the supported stack is in the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtm>. NetApp adds components on an ongoing basis. This online site is updated frequently to provide you with information about the current stack.

Software tools provided by the Host Utilities

The FCP AIX Host Utilities bundle the tools into a SAN Toolkit. The toolkit contains the following components:

- ◆ `aix_info` script. This script collects configuration information about your host. You use this script on AIX hosts that are not running VIO.
- ◆ `dotpaths` utility. This utility lets you set path priorities.
- ◆ `filer_info` script. This script collects information about the storage system.
- ◆ `ontap_collect`. This is snap extension that collects configuration information about your host just as the `aix_info` script does. You use `ontap_collect` in VIO server/client environments. VIO servers require you to run as the prime administrator (`padmin`) instead of `root`. You can execute `ontap_collect` when running as `padmin`.
- ◆ `san_version` command. This command displays the version of the SAN Toolkit that you are running.
- ◆ `sanlun` utility. You can use this utility to display information about LUNs on the storage system that are available to this host. In a VIO environment, you can only run the `sanlun` utility on the server; you cannot run it on a VIO client.
- ◆ Additional diagnostic scripts. The following scripts can also be used to collect troubleshooting information:
 - ❖ `brocade_info` collects information about Brocade switches installed in the network.
 - ❖ `cisco_info` collects information about Cisco switches installed in the network.
 - ❖ `mcdata_info` collects information about McData switches installed in the network.

- ❖ `qlogic_info` collects information about QLogic switches installed in the network.
- ◆ The man pages for `sanlun` and the diagnostic scripts.

Using a VIO environment with the Host Utilities

If you are using AIX with a VIO server, make sure you have your VIO environment set up before you install the Host Utilities. VIO environments enable you to share physical I/O resources between logical partitions. You can use both physical and virtual LUNs. You can create virtual disks on both the server and the client.

VIO servers require that you run as `padmin`, not `root`. As a result, some of the commands you use in setting up your Host Utilities environment are different than if you were running the Host Utilities on an AIX system without VIO. For example, to discover LUNs on VIO servers, you use the `cfgdev` command; on AIX hosts without VIO, you use the `cfgmgr` command.

Some other differences include:

- ◆ Installing the Host Utilities in VIO environments:
 - ❖ You must login as `padmin` and then use the `oem_setup_env` command to change to `root`. The Host Utilities require you to be `root` during the installation.

Note

When you enter the `oem_setup_env` command, the system goes into the `oem_setup_env` command mode. You can do things like install the Host Utilities, run the NetApp `dotpaths` utility, or run the `sanlun` utility when you are in this mode. To return to the normal operating mode for VIO environments, halt the `oem_setup_env` command mode. For example, enter `exit` at the command prompt.

- ❖ You must install the SAN Toolkit software package on the VIO server. In addition, you should also install it on each VIO client. While you cannot run the `sanlun` utility on the client, you can run the diagnostic scripts that are installed as part of the SAN Toolkit.
- ◆ Troubleshooting on VIO servers
 - ❖ To collect information about the host, you must run the `ontap_collect` file, not the `aix_info` file.

- ❖ If you have virtual SCSI LUNs set up, the `sanlun` utility does not recognize them. This is because the LUNs are identified as VID “IBM” and PID “VDASD.”

Note

For details on using VIO servers, please see the IBM documentation. If you want to see which commands you can use with VIO servers, enter the `help` command. Then use the `man` command to see details on each command.

About the HBAs

Do not use different HBA brands on a single host. The supported HBAs should be installed before you install the Host Utilities.

For information on the supported HBAs, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Multipathing

The FCP AIX Host Utilities for Native OS use AIX MPIO to provide multipathing. MPIO is AIX's native multipathing solution. By using MPIO you can configure multiple network paths between the host and storage system. If one path fails, FCP traffic continues on the remaining paths. MPIO is required if the host has multiple paths to a LUN.

The SAN Toolkit `sanlun` command will show you all NetApp consolidated devices to which the host has access to.

ALUA support

The FCP AIX Host Utilities for Native OS require that you use storage systems that have Asymmetric Logical Unit Access (ALUA) enabled. Check the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> to see which versions of Data ONTAP support ALUA and are supported by the Host Utilities.

NetApp introduced ALUA support with Data ONTAP 7.2 and Single Image cfmode.

ALUA defines a standard set of SCSI commands for discovering and managing multiple paths to LUNs on Fibre Channel SANs. You still need MPIO to run on the host, but you do not need additional host-specific plug-ins.

Remove previous versions before installing Host Utilities

The FCP AIX Host Utilities for Native OS uses a different stack from previous versions of the Host Utilities. You should not have both products installed on the same host.

If you have an earlier version installed, you must remove it before you install this version. For more information, see “[Uninstalling the software packages](#)” on page 32.

FCP AIX Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

Installing the FCP AIX Host Utilities and setting up your system involves a number of steps that are performed on both the storage system and the host. The following sections provide information about the tasks you need to perform.

Note

Occasionally there are known problems that can affect your system setup. Please check the FCP AIX Host Utilities Release Notes to make sure there are no issues affecting your system setup.

Pre-requisites

Before you install the Host Utilities, you must perform the following tasks:

1. Set up your system:

- ❖ Host OS, technology levels, and appropriate updates
- ❖ HBAs
- ❖ Drivers
- ❖ Storage system and Data ONTAP
- ❖ Switches (if used)

See the *Compatibility and Configuration Guide for NetApp FCP and ISCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml> for current information about the system requirements.

2. Verify that your storage system is:

- ❖ Licensed correctly
- ❖ Using the correct cfmode
- ❖ Running the FCP service
- ❖ Configured to work with the target HBAs
- ❖ Set up to work with the AIX host and the initiator HBAs
- ❖ Set up with working volumes and qtrees (if desired)

Note

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

3. If you're using a fabric connection, verify that the switch is:
 - ❖ Set up correctly
 - ❖ Zoned (if your system uses switches)
 - ❖ Cabled according to the instructions that came with the switch
4. Verify that:
 - ❖ The power to the switches, disk shelves, storage systems, and hosts is turned on in the correct order
 - ❖ The host and the storage system can communicate

Installation overview

Installing the Host Utilities involves the following tasks:

1. Get a copy of the Host Utilities software packages.
 - ❖ Get the compressed file containing the host settings software. This file contains the Object Data Manager (ODM) files for NetApp LUNs.
 - ❖ Get the compressed file containing the SAN Toolkit software. This file contains scripts that you can run to collect information about your configuration.
You can download these files or copy them from the CD
 - ❖ Extract the software packages from the compressed files.
2. Install the software packages:
 - ❖ Login as root. If you are on a VIO server, login as padm and then use the `oem_setup_env` command to change to root.
 - ❖ Go to the directory containing the extracted software package.
 - ❖ Use either the AIX System Management Interface Tool (SMIT) or the `installp` command to install the host settings software and the toolkit software.

Note

If you are using a VIO environment, you must install the SAN Toolkit software on the server and you should install it on each client. You do not need to install the host settings on the clients.

3. Verify that the HBA initiator queue depth is appropriate for your system.
 - ❖ Use the `lsattr -El` command to check the queue depth.
4. (Recommended) Check your cluster failover mode (cfmode).

- ❖ If you have a storage cluster, migrate your `cfmode` setting to single image mode.
 - ❖ If you have a direct-attached storage, verify that the media type is correct.
5. Configure the LUNs by performing the following tasks:
- ❖ Create at least one igroup and at least one LUN.
One way to create igroups and LUNs is to use the `lun setup` command to create igroups and LUNs. Specify `aix` as the value for the `ostype` attribute.
 - ❖ Enable ALUA, if you haven't already done so.
 - ❖ Set the path priorities. If the storage system is configured for ALUA, it automatically handles the path priorities. If you are not using ALUA, use the `dotpaths` utility that comes with the Host Utilities to set the path priorities.
 - ❖ Map the LUNs to the igroup.
 - ❖ Configure the host to discover the LUNs.
 - ❖ Get the host to discover the new LUNs by using the `cfgmgr` command. If you are running in a VIO environment, use the `cfgdev` command.
 - ❖ Display information about the LUNs and HBA.
On an AIX host, you can use the `sanlun` command.

Note

If you are using virtual LUNs, the `sanlun` command does not list them. This is because the virtual SCSI LUNs are identified as VID “IBM” and PID “VDASD”, which the `sanlun` utility does not recognize.

Where to go for more information

Host and storage system information

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

If you need information about...	See the following product documentation...
Known issues, system requirements, and last minute updates	The <i>Release Notes</i> , which are available at http://now.netapp.com/NOW/knowledge/docs/san/
A summary of some of the commands you might use with the Host Utilities	The <i>Quick Reference</i> for the FCP AIX Host Utilities, which is available from the NOW site
The latest supported configurations	◆ <i>NetApp Interoperability Matrix Tool</i> at http://now.netapp.com/matrix/mtx/login.do ◆ <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
Supported Fibre Channel SAN topologies	<i>FCP Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
Configuring the storage system	◆ <i>Best Practices for Reliability: New System Installation</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/index.shtml ◆ <i>Data ONTAP® Software Setup Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Block Access Management Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml

If you need information about...	See the following product documentation...
Enabling ALUA	<i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml In particular, see the section “<i>Enabling ALUA for a Fibre Channel igroup.</i>”
Managing SAN storage on the storage system	◆ <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Data ONTAP® Release Notes</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>Command Reference</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml ◆ <i>FilerView® Online Help</i> - available from within the FilerView GUI ◆ <i>Best Practices for Reliability: Ongoing Maintenance</i> at http://now.netapp.com/NOW/knowledge/docs/other/best_practices/newsys/index.htm
Upgrading Data ONTAP	◆ <i>Data ONTAP® Upgrade Guide</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml
Verifying compatibility of storage system with environmental requirements	◆ <i>Site Requirements</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Installing and configuring the HBA in your host	◆ Refer to your HBA vendor documentation.
Your host operating system	◆ Refer to your operating system documentation.
Working with VIO environments	◆ Refer to your operating system documentation. Additional documentation on using VIO environments is available on the IBM website.

About this chapter This chapter describes procedures for installing the FCP AIX Host Utilities.

**ATopics in this
chapter**

This chapter discusses the following topics:

- ◆ [“Completing the pre-installation tasks”](#) on page 14
- ◆ [“Obtaining the software packages”](#) on page 15
- ◆ [“Installing the Host Utilities software packages”](#) on page 18
- ◆ [“Uninstalling the software packages”](#) on page 32

Completing the pre-installation tasks

Before you begin

The FCP AIX Host Utilities for Native OS provides support for the AIX native stack and for VIO server. Your host must be running the supported HBAs and MPIO.

If you are using an environment, such as VIO, you must set up that environment also.

Before you install the SAN Toolkit software package:

- ◆ You must verify that your host system meets requirements before you begin the installation. Refer to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/ for the latest hardware and software requirements.
- ◆ If you currently have the FCP AIX Host Utilities 4.0 or FCP AIX Attach Kit installed, you must remove it. See and “[Uninstalling the software packages](#)” on page 32 or “[Uninstalling the software packages](#)” on page 32.

Once your host is set up, you must install the host settings software package and the SAN Toolkit software package.

For more information on planning your installation, see “[Confirming that your system is correctly set up](#)” on page 100.

Obtaining the software packages

Software packages

The FCP AIX Host Utilities include:

- ◆ The host settings file. This is the software package for the FCP AIX Host Utilities for Native OS using MPIO that contains NetApp's FCP device definitions, which are required for FCP AIX configurations. This is an Object Data Manager (ODM) stanza LPP package.
- ◆ The NetApp SAN Toolkit file. This software package contains the `sanlun` utility and diagnostic scripts. The `sanlun` utility displays configuration information about LUNs and HBAs. The diagnostic scripts collect information about your system if a problem occurs. Customer support may ask you to run the `sanlun` utility or a diagnostic script if a problem occurs.

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

Note

You use these same files for hosts running AIX and VIO.

Obtaining the software packages

You can get these files from the CD-ROM in your Host Utilities media kit. See “[Uncompressing the software files](#)” on page 17.

You can also download the files from the NOW web site at <http://now.netapp.com/NOW/cgi-bin/software>. You must download both the host settings file and the SAN Toolkit file. For information about downloading these files, see “[Downloading the software from the NOW site](#)” on page 16.

Once you have the files, you must uncompress them to get the software packages. To install the software, make sure the software packages are in a working directory on your AIX host.

Getting the files from the CD-ROM

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

Step	Action
1	Insert the CD-ROM into the CD-ROM drive.

Step	Action
2	Change to the root directory of the CD-ROM: # cd <CDROM_mountpoint>
3	Copy the host settings software file and the SAN Toolkit package from the CD-ROM's root directory to a temporary location on the host. Example: If you mount the CD-ROM on /mnt and want to place the SAN Toolkit file in the /tmp directory, you might enter commands similar to the following: # mount -v cdrfs -o ro /dev/cd0 /mnt # cp /mnt/santoolkit_aix_3.1.tar.Z /tmp # cd /tmp

Downloading the software from the NOW site

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 Host Utilities (Attach kits) 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.

Step	Action
5	Download the NetApp host settings software file and the SAN Toolkit software file to a local directory on your host. You must download both files. Note _____ If you download the files to a machine other than the AIX host, you must move them to the host.

Uncompressing the software files To uncompress the software files, complete the following steps.

Step	Action
1	Change to the directory where you saved the downloaded software and enter the following commands to uncompress the FCP AIX host settings package and the SAN Toolkit package: <pre># zcat ntap_aix_mpio_4.1.tar.Z tar -xvf - # zcat santoolkit_aix_3.1.tar.Z tar xvf -</pre> Result: The software packages are extracted. Example: Executing the <code>zcat</code> and <code>tar</code> commands with the host settings file and the SAN Toolkit file produces output similar to the following. <pre>HostA# zcat ntap_aix_mpio_4.1.tar.Z tar xvf - # zcat ntap_aix_mpio_4.1.tar.Z tar xvf - x ./ntap_aix_mpio_4.1 /NetApp.mpio_attach_kit, 195584 bytes, 382 media blocks. HostA# zcat santoolkit_aixv3.1.tar.Z tar xvf - x ./NetApp.SAN_toolkit, 1146880 bytes, 2240 media blocks. HostA#</pre>
2	Go to “ Installing the Host Utilities software packages ” on page 18.

Installing the Host Utilities software packages

Methods for installing the software

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

Note

This section uses SMIT as its installation example.

Upgrading from an earlier attach kit

If you have an FCP AIX Host Attach Kit or a previous version of the Host Utilities installed, you should remove it before you install a new version of the Host Utilities. See “[Uninstalling the software packages](#)” on page 32.

Note

Previous versions of the Host Utilities were called Attach Kits.

Using SMIT to install the software

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

Note

This installation example installs the host settings software first. You then repeat the steps to install the SAN Toolkit. If you have a VIO environment, you must install the SAN Toolkit on each client.

Step	Action
1	Make sure you are logged in as root. If you are running VIO on your host and you have logged in as padmin, enter the <code>oem_setup_env</code> command to become root. Once you execute this command you are in <code>oem_setup_env</code> command mode until you stop the mode by entering a command such as <code>exit</code>. Example: The following is a example of becoming root on a VIO server. <pre>vioserver\$ oem_setup_env</pre>
2	Start SMIT by entering the following command: <pre># smit install</pre>

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

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

Step	Action
5	At the Install Software screen, specify the location of the software in one of the following ways: ◆ Manually enter the location by providing the following information. ❖ If you're installing from the CD-ROM, enter the CD-ROM drive. ❖ If you're installing from the host machine, enter the path to the software package (for example, /var/tmp/ntap_aix_mpio_4.1/NetApp.mpio_attach_kit). ◆ Press F4 to display a list of options. If you want to use the F4 method, complete the following steps: 1. Press F4. 2. At the prompt for software to install, enter NetApp.mpio_attach_kit. 3. At the prompt asking whether you want to continue or cancel, press Enter to complete the installation. Example 1: The following is an example of entering the path to the software package when you're at the Install Software screen. <pre> +-----+ Install Software Type or select a value for the entry field. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software [/var/tmp/ntap_aix_mpio_4.1/NetApp.mpio_attach_kit]+ F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
	Example 2: After you enter path to the software package, SMIT displays the following screen. This is the screen where you enter name of the software package, Netapp.mpio_attach_kit, and press Enter. <pre> +-----+ Install Software Type or select values in entry fields. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software [/var/tmp/ntap_aix_mpio_4.1/NetApp.mpio_attach_kit] * SOFTWARE to install [NetApp.mpio_attach_kit] PREVIEW only? (install operation will NOT occur) no COMMIT software updates? yes SAVE replaced files? no AUTOMATICALLY install requisite software? yes EXTEND file systems if space needed? yes OVERWRITE same or newer versions? no VERIFY install and check file sizes? no Include corresponding LANGUAGE filesets? yes DETAILED output? no Process multiple volumes? yes ACCEPT new license agreements? no Preview new LICENSE agreements? no F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

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

Step	Action
7	SMIT displays information about the installation. Note_____ By default, the host settings software installs the filesets for both FCP and iSCSI. This is not a problem. If you are using only FCP, ignore the iSCSI fileset. _____ Example of a successful installation: The following is an example of output you might see when an installation successfully completes. <pre> COMMAND STATUS Command: running stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.mpio_attach_kit.pcmadm 4.1.0.0 NetApp.mpio_attach_kit.iscsi 4.1.0.0 NetApp.mpio_attach_kit.fcp 4.1.0.0 NetApp.mpio_attach_kit.config 4.1.0.0 << Copyright notice for NetApp.mpio_attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2007. All rights reserved. << End of copyright notice for NetApp.mpio_attach_kit >>. . . . Finished processing all filesets. (Total time: 5 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23825 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. * * * A T T E N T I O N * * * System boot image has been updated. You should reboot the system as soon as possible to properly integrate the changes and to avoid disruption of current functionality.</pre>

Step	Action
	<pre> installp: bosboot process completed. +-----+ Summaries: +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.mpio_attach_kit.pcmmod 4.1.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.iscsi 4.1.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.fcp 4.1.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.conf 4.1.0.0 USR APPLY SUCCESS COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.mpio_attach_kit.config 4.1.0.0 << Copyright notice for NetApp.mpio_attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2007. All rights reserved. << End of copyright notice for NetApp.mpio_attach_kit >>. . . . Finished processing all filesets. (Total time: 5 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23825 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. installp: bosboot process completed. +-----+ Summaries: +-----+ Installation Summary ----- </pre>

To install the SAN Toolkit software, enter the command line:

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

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

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

Selected Filesets
-----
NetApp.SAN_toolkit.sanlun 3.1.0.0           # Network Appliance SAN Toolki...
NetApp.SAN_toolkit.scripts 3.1.0.0         # Network Appliance SAN Toolki...

<< End of Success Section >>

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

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

installp:  APPLYING software for:
           NetApp.SAN_toolkit.scripts 3.1.0.0
           NetApp.SAN_toolkit.sanlun 3.1.0.0
. . . . . << Copyright notice for NetApp.SAN_toolkit >> . . . . .
(C) Copyright Network Appliance, Inc. 2003-2007.
All rights reserved.
. . . . . << End of copyright notice for NetApp.SAN_toolkit >>. . . .
Finished processing all filesets.  (Total time:  2 secs).
```

Summaries:				
Installation Summary				
Name	Level	Part	Event	Result
NetApp.SAN_toolkit.scripts	3.1.0.0	USR	APPLY	SUCCESS
NetApp.SAN_toolkit.sanlun	3.1.0.0	USR	APPLY	SUCCESS
#				

Sample installation using installp on a VIO server

The following is an example of using the `installp` command to install the Host Utilities on a VIO server.

```
vioserver$ oem_setup_env
# PS1="$ (hostname)# "
vioserver# ls -ls
total 1080
 184 -rw-r--r--   1 root      staff          186843 Nov 02 14:30 ntap_aix_mpio_4.1.tar.Z
 896 -rw-r--r--   1 root      staff          910519 Nov 02 14:30 santoolkit_aix_3.1.tar.Z
vioserver# zcat ntap_aix_mpio_4.1.tar.Z | tar xvf -
x ./ntap_aix_mpio_4.1/NetApp.mpio_attach_kit, 195584 bytes, 382 media blocks.
vioserver# cd ntap_aix_mpio_4.1
vioserver# installp -aXd NetApp.mpio_attach_kit NetApp.mpio_attach_kit
+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

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

```
Selected Filesets
-----
NetApp.mpio_attach_kit.config 4.1.0.0      # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.fcp 4.1.0.0        # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.iscsi 4.1.0.0       # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.pcmoem 4.1.0.0     # Network Appliance MPIO PCM K...

<< End of Success Section >>

FILESET STATISTICS
-----
    4 Selected to be installed, of which:
      4 Passed pre-installation verification
----
    4 Total to be installed

0503-409 installp: bosboot verification starting...
installp: bosboot verification completed.
+-----+
+-----+ Installing Software... +-----+
+-----+
installp: APPLYING software for:
          NetApp.mpio_attach_kit.pcmoem 4.1.0.0
          NetApp.mpio_attach_kit.iscsi 4.1.0.0
          NetApp.mpio_attach_kit.fcp 4.1.0.0
          NetApp.mpio_attach_kit.config 4.1.0.0

. . . . . << Copyright notice for NetApp.mpio_attach_kit >> . . . . .
(C) Copyright Network Appliance, Inc. 2003-2007.
All rights reserved.

. . . . . << End of copyright notice for NetApp.mpio_attach_kit >>. . . .

Finished processing all filesets. (Total time: 6 secs).

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

bosboot: Boot image is 30522 512 byte blocks.
0503-292 This update will not fully take effect until after a
system reboot.
```


Uninstalling the software packages

Before you uninstall previous versions of software

- Before you uninstall a previous version of the software, make sure you do the following;
- ◆ Unmount the file systems that contain the volume groups mapped to storage system LUNs.
 - ◆ Quiesce and export the volume groups.
 - ◆ Stop multipathing software.
 - ◆ Remove the existing NetApp LUN handles.
- These procedures disconnect NetApp LUNs from the host before you uninstall the software.

Uninstalling the software from an AIX host

You can use SMIT to uninstall the Host Utilities software.

Note_____

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

Uninstalling the software using SMIT

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

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

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

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

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

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

Step	Action
	Example of successful uninstall: The following output appears when you successfully uninstall the software. <pre>Command: running stdout: yes stderr: no Before command completion, additional instructions may appear below. 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. +-----+ Deinstalling Software... +-----+ installp: DEINSTALLING software for: NetApp.mpio_attach_kit.config 4.1.0.0 NetApp.mpio_attach_kit.fcp 4.1.0.0 NetApp.mpio_attach_kit.iscsi 4.1.0.0 NetApp.mpio_attach_kit.pcmadm 4.1.0.0 Finished processing all filesets. (Total time: 4 secs). 0503-409 installp: bosboot verification starting... installp: bosboot verification completed. 0503-408 installp: bosboot process starting... bosboot: Boot image is 23818 512 byte blocks. 0503-292 This update will not fully take effect until after a system reboot. * * * A T T E N T I O N * * * System boot image has been updated. You should reboot the system as soon as possible to properly integrate the changes and to avoid disruption of current functionality. installp: bosboot process completed. +-----+ Summaries: +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.mpio_attach_kit.confi 4.1.0.0 USR DEINSTALL SUCCESS NetApp.mpio_attach_kit.fcp 4.1.0.0 USR DEINSTALL SUCCESS NetApp.mpio_attach_kit.iscsi 4.1.0.0 USR DEINSTALL SUCCESS NetApp.mpio_attach_kit.pcmad 4.1.0.0 USR DEINSTALL SUCCESS</pre>

About this chapter

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

Topics covered in this chapter

This chapter includes the following topics:

- ◆ [“Deciding whether to migrate to single image mode”](#) on page 40
- ◆ [“Understanding `cfmode`”](#) on page 41
- ◆ [“Tasks to complete before you change the `cfmode`”](#) on page 45
- ◆ [“Upgrading to single image mode”](#) on page 47

Deciding whether to migrate to single image mode

When to migrate to single image mode

If you upgraded to Data ONTAP 7.1 or higher, it is recommended to change your storage systems' 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:

```
fcg 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 41.

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 4, “[Creating, Configuring, and Managing LUNs](#),” on page 53.

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 AIX hosts and you want to add HP-UX or Solaris 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.

Single image mode works with all currently supported cabling configurations. 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. The first figure is one example of cabling on a system with target HBAs that is using single image mode.

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 change the cfmode

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.
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 /filesystem_path
5	Quiesce the volume groups by entering the following command: varyoffvg volume_group_name
6	Export the volume group by entering the following command: exportvg volume_group_name
7	Enter the following command to list local devices: lsdev -Cc disk
8	Remove the existing NetApp LUN handles for the devices by entering the following command for each device: rmdev -dl device_name Example: This example removes existing LUN handles. # rmdev -dl hdisk2

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

Migrating the cluster to single image mode

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

Step	Action
1	Open a window to the console of each storage system you want to update.
2	Login at the console of the system(s) you are upgrading, and enter the following command to identify and resolve LUN and igroup mapping conflicts. <pre>lun config_check -S</pre> Result: Mappings that are displayed using the “*” character indicate an error, and also provide LUN and igroup information that you can use to resolve conflicts. Use the <code>lun map</code> command to re-map LUNS to a new ID that is unique to each node in the cluster. You use the <code>igroup</code> command to change igroup otypes, if required.
3	Switch to advanced privileges by entering the following command: <pre>priv set advanced</pre>
4	Stop the FCP service on both systems. On each system, enter: <pre>fcp 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 50.
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 cfmodes setting, make sure each HBA has one primary and at least one secondary path to each LUN. For additional information, see the FCP Configuration Guide at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/.

Reconfiguring AIX hosts

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

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

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

About this chapter

This chapter describes how to 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:

- ◆ [“Overview of completing the configuration”](#) on page 54
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 55
- ◆ [“LUN configuration overview”](#) on page 57
- ◆ [“Configuring LUNs with native AIX LVM”](#) on page 60
- ◆ [“Configuring LUNs with MPIO software on non-VIO hosts”](#) on page 62
- ◆ [“Configuring LUNs with MPIO software on VIO servers”](#) on page 66
- ◆ [“Creating volume groups and file systems on AIX hosts”](#) on page 73
- ◆ [“Using sanlun to display host LUN information”](#) on page 79
- ◆ [“Displaying host HBA information”](#) on page 82

Overview of completing the configuration

Task overview

Completing the configuration involves the following tasks.

Task	Discussion	Additional information
1. Create and map igroups and LUNs	An igroup is a collection of WWPNs on the storage system that map to one or more host HBAs. After you create the igroup, you must create LUNs on the storage system, and map the LUNs to the igroup.	For complete information, refer to your version of the <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
2. Enable ALUA	If you are using Data ONTAP 7.2 or later, you must have ALUA set up to work with igroups. To see if ALUA is set up for your igroup, use the <code>igroup show -v</code> command.	Refer to “ Check the ALUA settings on the storage system ” on page 103.
3. Discover LUNs	You can discover LUNs using either the <code>cfgmgr</code> command (AIX hosts without VIO) or the <code>cfgdev</code> command (hosts with VIO).	Refer to “ LUN configuration overview ” on page 57 for additional information.
4. (Optional) Create volume groups and file systems	You can create volume groups and file systems using the AIX tools that you normally use for these tasks.	Refer to “ Creating volume groups and file systems on AIX hosts ” on page 73 for additional information.
5. Display information about the LUNs and the HBAs	You can use the <code>sanlun</code> utility to display information about the LUNs and the HBAs.	Refer to “ Using sanlun to display host LUN information ” on page 79 and “ Displaying host HBA information ” on page 82 for additional information.

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.

Note

If you have a cluster, you must create a separate igroup on each system in the cluster.

- ◆ 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.

Methods for creating igroups and LUNs

You can create igroups and LUNs on the storage system by:

- ◆ Entering the `lun setup` command
This method prompts you through the process of creating a LUN, creating an igroup, and mapping the LUN to the igroup.
- ◆ Entering a series of individual commands (such as `lun create`, `igroup create` and `lun map`).

Use this method to create one or more LUNs and igroups in any order.

For complete information about creating igroups and LUNs, refer to *Data ONTAP™ Block Access Management Guide for FCP* at

http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

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 60.
- ◆ If your configuration includes multipathing, refer to [“Configuring LUNs with MPIO software on non-VIO hosts”](#) on page 62
- ◆ If you have a VIO server, refer to [“Configuring LUNs with MPIO software on VIO servers”](#) on page 66.
- ◆ If you are unsure of your configuration and software, refer to [“LUN configuration overview”](#) on page 57.

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 MPIO.

Configuring LUNs with native AIX LVM (no multipathing)

To configure LUNs using native AIX LVM:

1. Use `cfgmgr` to get the host to discover the new LUNs.

Note

If you have a VIO server, you must use the `cfgdev` command instead of the `cfgmgr` command.

Refer to [“Configuring LUNs with native AIX LVM”](#) on page 60 for additional information.

Non-VIO hosts: configuring LUNs with native AIX LVM using MPIO

To configure LUNs using native AIX LVM with MPIO on hosts that are not running VIO:

1. Use `cfgmgr` to get the host to discover the new LUNs.

Note

If you have a VIO server, you must use the `cfgdev` command instead of the `cfgmgr` command.

2. After the command completes, use the `lsdev` command to verify that the `hdisk` MPIO devices were created.
3. Use the `lsattr -El hdisk2` command to check the system setup.
4. If you are using Data ONTAP 7.2 or later with ALUA enabled, ALUA automatically sets the path priorities. If you are not using ALUA, you must run the `dotpaths` utility on all the disks to order and prioritize the paths to the storage system.

5. Run the `sanlun lun show -p` command.

Refer to “[Configuring LUNs with MPIO software on non-VIO hosts](#)” on page 62 for additional information.

VIO servers: configuring LUNs with native AIX LVM using MPIO

While the steps for working with LUNs on VIO servers are similar to those on hosts using MPIO without VIO, some of the commands are different. This is because VIO servers use the `padmin` login instead of a root login.

To configure LUNs using native AIX LVM with MPIO on VIO servers, you need to change some of the commands you use:

1. Use `cfgdev` to get the host to discover the new LUNs.
2. After the command completes, use the `lsdev` command to verify that the `hdisk` MPIO devices were created.
3. Use the `lsdev -dev hdisk2 -attr` command to check the system setup.
4. Run the `oem_setup_env` command. This changes the mode to allow you to run NetApp utilities such as `dotpaths` and `sanlun`.
5. If you are using Data ONTAP 7.2 or later with ALUA enabled, ALUA automatically sets the path priorities. If you are not using ALUA, you must run the `dotpaths` utility on all the disks to order and prioritize the paths to the storage system.
6. Run the `sanlun lun show -p` command (you must run the `oem_setup_env` command before the `sanlun` command).
7. Halt the `oem_setup_env` command mode by entering a command such as `exit`.
8. Configure the new storage as desired. For example, you can create virtual disks, create a new volume group, add the `hdisk` to an existing volume group, create logical volume groups, or create file systems.

Refer to “[Configuring LUNs with MPIO software on VIO servers](#)” on page 66 for additional information.

Utility for displaying information

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

For more information, see:

- ◆ [“Using sanlun to display host LUN information”](#) on page 79
- ◆ [“Displaying host HBA information”](#) on page 82

Additional tasks

Once you’ve set up your LUNs, you may want to configure the new storage. For example, you can create a new volume group, add the hdisk to an existing volume group, create logical volume groups, or create file systems. You can also configure data access (raw, jfs, or jfs2 file system).

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>

Configuring the LUNs

Configuring LUNs with MPIO software on non-VIO hosts

Getting the host to discover the new LUNs

To get a host that is not running VIO 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 LUNs. ◆ Map the new LUNs as MPIO devices using the default PCM.
3	Execute <code>lsdev</code> to verify that the <code>hdisk</code> MPIO devices were created. Write down the <code>hdisk</code> instance numbers for path configuration. Example: The following is the type of output you would see when you run this command. <pre># lsdev -Ccdisk hdisk0 Available 08-08-00-5,0 16 Bit LVD SCSI Disk Drive hdisk1 Available 08-08-00-8,0 16 Bit LVD SCSI Disk Drive hdisk2 Available 04-08-02 MPIO NetApp FCP Default PCM Disk hdisk3 Available 04-08-02 MPIO NetApp FCP Default PCM Disk hdisk4 Available 04-08-02 MPIO NetApp FCP Default PCM Disk hdisk5 Available 04-08-02 MPIO NetApp FCP Default PCM Disk</pre>

Step	Action
4	Execute the <code>lsattr</code> command to get information about your setup. Example: The following is the type of output you would see when you run this command. <pre>lsattr -El hdisk65 PCM PCM/friend/NetAppDefaultPCM Path Control Module False PR_key_value Persistant Reserve Key Value True algorithm round_robin Algorithm True clr_q no Device CLEARS its Queue on error True dist_err_pcmt 0 Distributed Error Sample Time True dist_tw_width 50 Distributed Error Sample Time True hcheck_cmd inquiry Health Check Command True hcheck_interval 30 Health Check Interval True hcheck_mode nonactive Health Check Mode True location Location Label True lun_id 0x3f00000000000000 Logical Unit Number ID False lun_reset_spt yes LUN Level Reset True max_transfer 0x40000 Maximum TRANSFER Size True node_name 0x500a09808607cddc FC Node Name False pvid 00ce694b35ee0b26000000000000000 Physical volume identifier False q_err yes Use QERR bit True q_type simple Queuing TYPE True queue_depth 12 Queue DEPTH True reassign_to 120 REASSIGN time out value True reserve_policy no_reserve Reserve Policy True rw_timeout 30 READ/WRITE time out value True scsi_id 0xa0000 SCSI ID False start_timeout 60 START unit time out value True ww_name 0x500a09839607cddc FC World Wide Name False</pre>

Step	Action
5	Order and prioritize the paths to the storage system LUNs. ◆ If you are running a version of Data ONTAP that supports ALUA, ALUA automatically sets the paths. this. For information on using ALUA, see <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ If you are not running ALUA, you must execute the <code>dotpaths</code> utility to set the path priorities. The <code>dotpaths</code> utility is installed when you install the Host Utilities. If you enter <code>dotpaths</code> without any options, it sets the priority for all Netapp LUNs. This utility has the format: <code>dotpaths [-hqv] [hdiskN ...]</code> You can use the command line options to: ❖ Specify a set of disks for which you want to set the priority. When you enter a list of <code>hdisk</code> names, the utility only sets the priority for those disks. ❖ Display all the priorities. Enter the <code>-v</code> option to enable verbose mode. ❖ Query all or individual disk priorities. Enter the <code>-q</code> option to perform this task.

Step	Action																																																																																																						
6	Execute the <code>sanlun lun show -p</code> command to get information about your setup. Example: The following is the type of output you would see when you run this command. <pre># sanlun lun show -p ONTAP_PATH: shfas3020-07:/vol/lun7 LUN: 7 LUN Size: 2g (2147483648) Host Device: hdisk9 LUN State: GOOD Controller_CF_State: Cluster Enabled Controller Partner: shfas3020-08 Multipath Provider: AIX Native Multipathing Algorithm: round_robin</pre> <table><tr><th>MPIO path status</th><th>Controller path type</th><th>AIX MPIO path</th><th>host HBA</th><th>Controller target port</th><th>AIX MPIO path priority</th></tr><tr><td>Enabled</td><td>primary</td><td>path0</td><td>fcs0</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path1</td><td>fcs0</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path2</td><td>fcs0</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path3</td><td>fcs0</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path4</td><td>fcs1</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path5</td><td>fcs1</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path6</td><td>fcs1</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path7</td><td>fcs1</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path8</td><td>fcs2</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path9</td><td>fcs2</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path10</td><td>fcs2</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path11</td><td>fcs2</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path12</td><td>fcs3</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path13</td><td>fcs3</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path14</td><td>fcs3</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path15</td><td>fcs3</td><td>1b</td><td>1</td></tr></table>	MPIO path status	Controller path type	AIX MPIO path	host HBA	Controller target port	AIX MPIO path priority	Enabled	primary	path0	fcs0	1a	255	Enabled	secondary	path1	fcs0	1a	1	Enabled	primary	path2	fcs0	1b	255	Enabled	secondary	path3	fcs0	1b	1	Enabled	primary	path4	fcs1	1a	255	Enabled	secondary	path5	fcs1	1a	1	Enabled	primary	path6	fcs1	1b	255	Enabled	secondary	path7	fcs1	1b	1	Enabled	primary	path8	fcs2	1a	255	Enabled	secondary	path9	fcs2	1a	1	Enabled	primary	path10	fcs2	1b	255	Enabled	secondary	path11	fcs2	1b	1	Enabled	primary	path12	fcs3	1a	255	Enabled	secondary	path13	fcs3	1a	1	Enabled	primary	path14	fcs3	1b	255	Enabled	secondary	path15	fcs3	1b	1
MPIO path status	Controller path type	AIX MPIO path	host HBA	Controller target port	AIX MPIO path priority																																																																																																		
Enabled	primary	path0	fcs0	1a	255																																																																																																		
Enabled	secondary	path1	fcs0	1a	1																																																																																																		
Enabled	primary	path2	fcs0	1b	255																																																																																																		
Enabled	secondary	path3	fcs0	1b	1																																																																																																		
Enabled	primary	path4	fcs1	1a	255																																																																																																		
Enabled	secondary	path5	fcs1	1a	1																																																																																																		
Enabled	primary	path6	fcs1	1b	255																																																																																																		
Enabled	secondary	path7	fcs1	1b	1																																																																																																		
Enabled	primary	path8	fcs2	1a	255																																																																																																		
Enabled	secondary	path9	fcs2	1a	1																																																																																																		
Enabled	primary	path10	fcs2	1b	255																																																																																																		
Enabled	secondary	path11	fcs2	1b	1																																																																																																		
Enabled	primary	path12	fcs3	1a	255																																																																																																		
Enabled	secondary	path13	fcs3	1a	1																																																																																																		
Enabled	primary	path14	fcs3	1b	255																																																																																																		
Enabled	secondary	path15	fcs3	1b	1																																																																																																		

Determine which device to use as the physical volume

You must determine which device to use as the physical volume for the LVM before you create a volume group. You can use the `lsdev -Cc disk` command to determine which host disk device MPIO uses as the I/O access device for the LUN.

Configuring LUNs with MPIO software on VIO servers

Getting a VIO server to discover the new LUNs

To get a VIO server 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 to the host as padmin. Note_____ Some storage system commands, such as dotpaths and sanlun, are not recognized by padmin. To run these commands, you must go into oem_setup_env mode. Once you enter the oem_setup_env command, you stay in that mode until you halt it by using a command such as exit. .
2	Execute the cfgdev command on the AIX host to discover the newly created LUNs. Result: This command causes the VOI server to ◆ Log into the fabric and to check for new LUNs. ◆ Map the new LUNs as MPIO devices using the default PCM. Example: On a VIO server, use the cfgdev command to discover the LUNs. <pre>vioserver\$ cfgdev</pre>

Step	Action
3	Execute <code>lsdev</code> to verify that the <code>hdisk</code> <code>MPIO</code> devices were created. Write down the <code>hdisk</code> instance numbers for path configuration. Example: The following is the type of output you would see when you run this command. <pre> vioserver\$ lsdev -type disk name status description hdisk0 Available 16 Bit LVD SCSI Disk Drive hdisk1 Available 16 Bit LVD SCSI Disk Drive hdisk2 Available MPIO NetApp FCP Default PCM Disk hdisk3 Available MPIO NetApp FCP Default PCM Disk hdisk4 Available MPIO NetApp FCP Default PCM Disk hdisk5 Available MPIO NetApp FCP Default PCM Disk hdisk6 Available MPIO NetApp FCP Default PCM Disk hdisk7 Available MPIO NetApp FCP Default PCM Disk hdisk8 Available MPIO NetApp FCP Default PCM Disk hdisk9 Available MPIO NetApp FCP Default PCM Disk hdisk10 Available MPIO NetApp FCP Default PCM Disk hdisk11 Available MPIO NetApp FCP Default PCM Disk hdisk12 Available MPIO NetApp FCP Default PCM Disk hdisk13 Available MPIO NetApp FCP Default PCM Disk hdisk14 Available MPIO NetApp FCP Default PCM Disk hdisk15 Available MPIO NetApp FCP Default PCM Disk hdisk16 Available MPIO NetApp FCP Default PCM Disk hdisk17 Available MPIO NetApp FCP Default PCM Disk hdisk18 Available MPIO NetApp FCP Default PCM Disk hdisk19 Available MPIO NetApp FCP Default PCM Disk hdisk20 Available MPIO NetApp FCP Default PCM Disk hdisk21 Available MPIO NetApp FCP Default PCM Disk hdisk22 Available MPIO NetApp FCP Default PCM Disk hdisk23 Available MPIO NetApp FCP Default PCM Disk vioserver\$ </pre>

Step	Action
4	Now use lsdev to get information about your setup. Example: The following is the type of output you would see when you run this command. <pre> vioserver\$ lsdev -dev hdisk2 -attr attribute value description user_settable PCM PCM/friend/NetAppDefaultPCM Path Control Module False algorithm round_robin Algorithm True clr_q no Device CLEARS its Queue on error True dist_err_pcmt 0 Distributed Error Sample Time True dist_tw_width 50 Distributed Error Sample Time True hcheck_cmd inquiry Health Check Command True hcheck_interval 0 Health Check Interval True hcheck_mode nonactive Health Check Mode True location location Location Label True lun_id 0x200000000000000 Logical Unit Number ID False max_transfer 0x40000 Maximum TRANSFER Size True node_name 0x500a09808667f35b FC Node Name False pvid 00c0ca0e99e910ca00000000000000000 Physical volume identifier False q_err yes Use QERR bit True q_type simple Queuing TYPE True queue_depth 12 Queue DEPTH True reassign_to 120 REASSIGN time out value True reserve_policy no_reserve Reserve Policy True rw_timeout 30 READ/WRITE time out value True scsi_id 0xa0c00 SCSI ID False start_timeout 60 START unit time out value True ww_name 0x500a09859667f35b FC World Wide Name False vioserver\$ </pre>

Step	Action
5	Enter <code>oem_setup_env</code> mode so that you can execute the Host Utilities commands <code>dotpaths</code> (if necessary) and <code>sanlun</code>. To get into this mode, enter: <code>oem_setup_env</code> Note _____ Once you enter the <code>oem_setup_env</code> command, you stay in that mode until you halt it by using a command such as <code>exit</code>. _____ Example When you enter the <code>oem_setup_env</code> mode, the command line prompt changes from a dollar sign to a pound sign. <pre>vioserver\$ oem_setup_env # PS1="\$(hostname)# "</pre>

Step	Action
6	Order and prioritize the paths to the storage system LUNs. ◆ If you are running a version of Data ONTAP that supports ALUA, ALUA automatically sets the path priorities. For information on using ALUA, see <i>Data ONTAP® Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. ◆ If you are not running ALUA, you must execute the dotpaths utility to set the path priorities. The dotpaths utility is installed when you install the Host Utilities. If you enter dotpaths without any options, it sets the priority for all Netapp LUNs. This utility has the format: <pre>dotpaths [-hqv] [hdiskN ...]</pre> You can use the command line options to: ❖ Specify a set of disks for which you want to set the priority. When you enter a list of hdisk names, the utility only sets the priority for those disks. ❖ Display all the priorities. Enter the -v option to enable verbose mode. ❖ Query all or individual disk priorities. Enter the -q option to perform this task. The following is an example of using the dotpaths utility. <pre>vioserver# dotpaths -v hdisk2 (path 0): Optimized path - Setting priority to 255 hdisk2 (path 1): Optimized path - Setting priority to 255 hdisk2 (path 2): Unoptimized path - Current priority of 1 is correct hdisk2 (path 3): Unoptimized path - Current priority of 1 is correct ... Path priority set and/or verified for 22 disks, 86 total paths. vioserver#</pre>

Step	Action
7	Execute the <code>sanlun lun show -p</code> command to get information about your setup. Example: The following is the type of output you would see when you run this command. The output from this command has been truncated. This example also includes the <code>sanlun lun show -d hdisk9 -v</code> command to display information about <code>hdisk9</code>. <pre> vioserver# sanlun lun show -p ONTAP_PATH: shfas960-01:/vol/vol1/vioserver_7 LUN: 7 LUN Size: 250m (262144000) Host Device: hdisk9 LUN State: GOOD Controller_CF_State: Cluster Enabled Controller Partner: shfas960-02 Multipath Provider: AIX Native Multipathing Algorithm: round_robin ----- MPIO Controller AIX Controller AIX MPIO path path MPIO host target HBA path status type path HBA port priority ----- Enabled primary path0 fcs0 4a 255 Enabled primary path1 fcs0 4b 255 Enabled secondary path2 fcs0 4a 1 Enabled secondary path3 fcs0 4b 1 ... vioserver# sanlun lun show -d hdisk9 filer: lun-pathname device filename adapter protocol lun size lun state shfas960-01: /vol/vol1/vioserver_7 hdisk9 fcs0 FCP 250m (262144000) GOOD Serial number: VrHqko8bg72t Filer FCP nodename:500a098083e152d9 Filer FCP portname:500a098193e152d9 Filer adapter name: v.4a Filer IP address: 10.60.240.66 Filer volume name:vol1 FSID:0xe10152d9 Filer qtree name:/vol/vol1 ID:0x0 Filer snapshot name: ID:0x0 vioserver# </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 the `lsdev` command to display information about the available disks.

Example: The following example shows part of the output of the `lsdev-type disk` command.

```
vioserver$  
$ lsdev -type disk  
hdisk0          Available  MPIIO NetApp FCP Default PCM Disk  
hdisk1          Available  MPIIO NetApp FCP Default PCM Disk  
hdisk2          Available  MPIIO NetApp FCP Default PCM Disk  
hdisk3          Available  MPIIO NetApp FCP Default PCM Disk  
hdisk4          Available  MPIIO NetApp FCP Default PCM Disk  
vioserver$
```

Creating volume groups and file systems on AIX hosts

Additional tasks for working with LUNs

You can also set up volume groups and file systems.

Note_____

This section applies to AIX hosts only.

Creating a volume group

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

Note_____

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

Step	Action
1	Once you have determined which device to use as the physical volume for the volume manager, you can use the AIX System Management Interface Tool (SMIT) to create a volume group. Use the <code>smit vg</code> command to start the tool.

Step	Action																																										
2	Enter a volume group name and the information required 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">-----</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	-----		
	[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																																									

3	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. 0516-1254 mkvg: Changing the PVID in the ODM. newvg <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr><tr><td>n=Find Next</td><td></td><td></td><td></td></tr><tr><td><F10 exits smit></td><td></td><td></td><td></td></tr><tr><td colspan="4">-----</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 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 “Configuring the LUNs” on page 62. ----- Add a Standard Journaled File System Type or select values in entry fields. Press Enter AFTER making all desired changes. [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 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									

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

Step	Action
	So you might enter: ◆ <code>sanlun lun show -p</code> to display a listing of all the paths associated with the LUN. This information is useful if you need to set up path ordering or troubleshoot a problem with path ordering. ◆ <code>sanlun lun show -d /dev/hdisk<x></code> to display the summary listing of the LUN(s) associated with the host device <code>/dev/hdisk<x></code> (<i>x</i> is the number of the hdisk; for example, you might have <code>hdisk1</code> or <code>hdisk4</code>). ◆ <code>sanlun lun show -v all</code> to display verbose output for all the LUN(s) currently available on the host. ◆ <code>sanlun lun show toaster</code> to display a summary listing of all the LUNs available to the host served by the storage system called <code>toaster</code>. ◆ <code>sanlun lun show toaster:/vol/vol0/lun0</code> to display a summary listing of all the LUNs available to the host served by <code>lun0</code> on <code>toaster</code>. Note_____ When you specify either the <code>sanlun lun show <storage_syatem_name></code> or the <code>sanlun lun show <storage_syatem_name:storage_syatem_pathname></code> command, the utility displays only the LUNs that have been discovered by the host. LUNs that have not been discovered by the host are not displayed. _____

Step	Action																																																																																																						
	Example 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 MPIO 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 ONTAP_PATH: filerA:/vol lun7 LUN: 7 LUN Size: 2g (2147483648) Host Device: hdisk9 LUN State: GOOD Controller_CF_State: Cluster Enabled Controller Partner: filerB Multipath Provider: AIX Native Multipathing Algorithm: round_robin</pre> <table><tr><th>MPIO path status</th><th>Controller path type</th><th>AIX MPIO path</th><th>host HBA</th><th>Controller target port</th><th>AIX MPIO priority</th></tr><tr><td>Enabled</td><td>primary</td><td>path0</td><td>fcs0</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path1</td><td>fcs0</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path2</td><td>fcs0</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path3</td><td>fcs0</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path4</td><td>fcs1</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path5</td><td>fcs1</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path6</td><td>fcs1</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path7</td><td>fcs1</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path8</td><td>fcs2</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path9</td><td>fcs2</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path10</td><td>fcs2</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path11</td><td>fcs2</td><td>1b</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path12</td><td>fcs3</td><td>1a</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path13</td><td>fcs3</td><td>1a</td><td>1</td></tr><tr><td>Enabled</td><td>primary</td><td>path14</td><td>fcs3</td><td>1b</td><td>255</td></tr><tr><td>Enabled</td><td>secondary</td><td>path15</td><td>fcs3</td><td>1b</td><td>1</td></tr></table>	MPIO path status	Controller path type	AIX MPIO path	host HBA	Controller target port	AIX MPIO priority	Enabled	primary	path0	fcs0	1a	255	Enabled	secondary	path1	fcs0	1a	1	Enabled	primary	path2	fcs0	1b	255	Enabled	secondary	path3	fcs0	1b	1	Enabled	primary	path4	fcs1	1a	255	Enabled	secondary	path5	fcs1	1a	1	Enabled	primary	path6	fcs1	1b	255	Enabled	secondary	path7	fcs1	1b	1	Enabled	primary	path8	fcs2	1a	255	Enabled	secondary	path9	fcs2	1a	1	Enabled	primary	path10	fcs2	1b	255	Enabled	secondary	path11	fcs2	1b	1	Enabled	primary	path12	fcs3	1a	255	Enabled	secondary	path13	fcs3	1a	1	Enabled	primary	path14	fcs3	1b	255	Enabled	secondary	path15	fcs3	1b	1
MPIO path status	Controller path type	AIX MPIO path	host HBA	Controller target port	AIX MPIO priority																																																																																																		
Enabled	primary	path0	fcs0	1a	255																																																																																																		
Enabled	secondary	path1	fcs0	1a	1																																																																																																		
Enabled	primary	path2	fcs0	1b	255																																																																																																		
Enabled	secondary	path3	fcs0	1b	1																																																																																																		
Enabled	primary	path4	fcs1	1a	255																																																																																																		
Enabled	secondary	path5	fcs1	1a	1																																																																																																		
Enabled	primary	path6	fcs1	1b	255																																																																																																		
Enabled	secondary	path7	fcs1	1b	1																																																																																																		
Enabled	primary	path8	fcs2	1a	255																																																																																																		
Enabled	secondary	path9	fcs2	1a	1																																																																																																		
Enabled	primary	path10	fcs2	1b	255																																																																																																		
Enabled	secondary	path11	fcs2	1b	1																																																																																																		
Enabled	primary	path12	fcs3	1a	255																																																																																																		
Enabled	secondary	path13	fcs3	1a	1																																																																																																		
Enabled	primary	path14	fcs3	1b	255																																																																																																		
Enabled	secondary	path15	fcs3	1b	1																																																																																																		

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

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 86
- ◆ “[Setting up a SAN Boot LUN](#)” on page 89
- ◆ “[Cloning the root volume group](#)” on page 94

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.

Contents of this chapter

This chapter provides an example of setting up a boot LUN. The examples uses a host running the AIX B 5L. The procedure described in this chapter is just one example of creating a bootable LUN. There are other methods. You can apply this procedure to most configurations.

This chapter does not describe other examples of bootable LUNs, such as using the AIX Alternate Disk Install, which allows you to install the OS while it is up and running.

Refer to AIX documentation for details about additional configuration methods.

Important notes about creating a SAN boot LUN

Before you start make sure the following is true:

- ◆ You are running the correct version and technology level of the operating system. See the online Release Notes for more information.
- ◆ You are logged in as root.

When you create the SAN boot LUN, make sure you set it so that you do not have more than four paths to the boot LUN.

Note

If you are not using a FAS270, FAS3000, or FAS6000 storage system, you must limit the number of paths to one.

Pre-requisites

The examples in this chapter assume that the host and storage system are configured properly and use software and firmware that is supported by the FCP AIX Host Utilities.

Configuration overview

Configuring SAN booting involves:

1. Configuring the igroup and SAN boot LUN.
 - a. Determine the size of the boot LUN and creating the igroup.
 - b. Create the boot LUN.
 - c. Map the boot LUN to the igroup.
2. Setting up your system
 - a. Create a portset for at least one target port and a failover port to an igroup that will only be used for the boot LUN.

Note

For FAS270, FAS3000, or FAS6000 storage systems, limit the number of paths to the LUN to four. For other storage systems, limit the number of paths to one.

- b. Bind the portset to an igroup that is used **only** for the boot LUN.
3. Installing the OS on the boot LUN.
 - a. Install the operating system and reboot.
 - b. Install the he Host Utilities software on the boot LUN and reboot.
 - c. Remove the SCSI-2 reservation placed on the boot LUN from the initial installation by halting the operating system and executing the online and offline command on the LUN.

Note

If you are using an older storage system you can add ports back to the portset to bring it up to the four-path maximum.

- d. Power on the server and boot from the LUN.
 - e. Use the `rmdev` command to clean up any left over ghost devices.
 - f. Set the path priorities. ALUA automatically handles this. If you are not using ALUA, run the `dotpaths` utility to set the path priorities.

- g.** Invalidate the current bootlist and reset it to use the primary boot devices.

Setting up a SAN Boot LUN

Important information about creating boot LUNs

To create a boot LUN requires the following:

- ◆ The correct operating system and technology levels.
- ◆ A limit of four paths to the boot LUN. Your data LUNs can use all four paths to the storage array,

Creating a SAN boot LUN

Creating a LUN you can boot from involves several tasks that you must perform on the host and on the storage system. The following steps tell you how to create the boot LUN.

Note_____The steps for installing the operating system are ones you would use if you were installing from a CD-ROM. For information on cloning the operating system, see [“Cloning the root volume group”](#) on page 94.

Step	Action
1	Determine the size of the boot LUN. You must create a LUN that is large enough to provide space for the AIX operating system version you are using. See the AIX product documentation to determine how much disk space is required.

Step	Action
2	Create an igroup. You must be at your host logged in as root to create an igroup. You perform some steps on the host (Host) and some steps on the storage system (Storage system). You can use a second session on the host to telnet to the storage system, if you choose. (Host) Get the World Wide Port Name (WWPN) for each of the host's HBAs. You can get this off of the adapter or, if the host is powered up, by running local disk commands or the <code>sanlun fcp show adapter -c</code> command. - Run the <code>igroup create</code> command to get information necessary to create the igroup. Example: If you had an AIX system called boothost with two HBAs, this command produce the following output on the console. <pre>igroup create -f -t aix "boothost" 100a0b000010a700 100a0b000010a606</pre> (Storage system) Copy the output from the <code>igroup create</code> command and paste it onto the command line of the storage system you telnetted to. (Storage system) Press enter. The storage system creates an igroup using the name of the AIX host. For example, the command line in step 1 creates an igroup called boothost. (Storage system) Use the <code>igroup show</code> command line to verify that the igroup was created. Example: If you had set up an igroup based on the command line in step 1, the <code>igroup show</code> command would produce the following result. <pre>filerA> igroup show -v boothost boothost (FCP) (ostype: aix): 10:00:00:00:c9:40:56:23 (logged in on: vtic, 1a, 0c) 10:00:00:00:c9:40:56:e9 (logged in on: vtic, 1b, 0d)</pre> For details on creating igroups, see the <i>Data ONTAP™ Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
3	Create the boot LUN and map it to an igroup. You can use the <code>lun setup</code> command to perform these tasks. This command steps you through the process. For complete information, see the <i>Data ONTAP™ Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
4	Use the <code>lun show</code> command to confirm that the LUN you created is mapped to the igroup. <pre>filerA> lun show -g boothost /vol/sanboot/boothost 50g (53687091200) (r/w, online, mapped)</pre>

Step	Action
5	Create a portset that has at least one target port and one failover port to the igroup that the boot LUN uses. This igroup can only be used for the boot LUN. You must create this portset on the storage system where you created the boot LUN. Use the command: <pre>portset create -f <PortSetName> <FilerTargetPortName></pre> <i>PortSetName</i> is the name of the portset you are creating. It can be up to 95 characters long. <i>FilerTargetPortName</i> is the name of the FCP target port. Example: This example creates a portset called <code>sanboot</code> for the LUN. <pre>filerA> portset create -f sanboot</pre>
6	Bind the portset to the igroup for the boot LUN. Use the command: <pre>igroup bind <iGroupName> <PortSetName></pre> <i>iGroupName</i> is the name of the igroup you created; for example, <code>boothost</code>. <i>PortSetName</i> is the name of the portset you created; for example, <code>sanboot</code>. Example: This example creates a portset called <code>sanboot</code> for the LUN associated with target port <code>0c</code> on <code>filerA</code>. It then uses the <code>portset show</code> command to confirm the port has been set up. <pre>filerA> igroup bind boothost sanboot filerA> portset show sanboot (FCP): ports: filerA 0c igroups: boothost filerA> igroup show boothost boothost2 (FCP) (ostype: aix): 10:00:00:00:c9:50:54:5a (logged in on: vtic, 0c) 10:00:00:00:c9:50:5f:8c (logged in on: 1b, vtic, 0d) Pset: sanboot</pre>
7	Install the AIX operating system on the boot LUN and reboot. For details on installing the AIX operating system, follow the instructions in the AIX documentation.
8	Install the Host Utilities. For information on doing this, see “Installing the Host Utilities software packages” on page 18.

Step	Action
9	Halt the server. You can use the <code>shutdown now</code> command.
10	Remove the SCSI-2 reservation that the initial installation placed on the boot LUN. Use the <code>lun offline</code> and the <code>lun online</code> commands. Note _____ If you are not using a FAS270, FAS3000, or FAS6000 storage system, you can add ports to the portset to bring the number of paths to the LUN up to a maximum of four. For information on working with portsets, see the <i>Data ONTAP™ Block Access Management Guide for FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. _____ Example: This example removes the SCSI-2 reservation. <pre>filerA> lun offline /vol/sanboot/boothost filerA> lun online /vol/sanboot/boothost</pre>
11	Restart the server so that it boots from the boot LUN.
12	Use the <code>sanlun lun show -p</code> command to verify that the LUN only has four paths.
13	Use the <code>rmdev</code> command to clean up any left over ghost devices that resulted from the ODM device overlay from MPIO to MPIO NetApp. Then use the <code>cfgmgr</code> command to rescan the bus for MPIO NetApp Fibre Channel devices.
14	If you are not using ALUA, set the path priorities by running the <code>dotpaths</code> utility that comes with the Host Utilities. Example: This example runs the <code>dotpaths</code> utility in verbose mode. <pre>host> /usr/sbin/dotpaths -v</pre>

Step	Action
15	Invalidate the current bootlist and create a bootlist that boots only from the primary paths of the boot LUN. You can use the following commands to perform these steps: <pre>bootinfo -b bootlist -m normal -i bootlist -m normal <BootDeviceHandle></pre> These commands: ◆ Return information on the last boot device. ◆ Invalidate the normal mode boot list ◆ Set a the new boot device to a normal boot list. Note _____ This procedure helps prevent a common boot issue where the bootlist becomes filled with paths that cannot be used during the boot cycle. _____

Cloning the root volume group

Cloning the operating system

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

Step	Action
1	Enter the <code>alt_disk_install</code> command and specify the root volume that you want to clone. Example: The following command clones <code>hdisk1</code>. <pre>-> alt_disk_install -C hdisk1 +-----+ --+ ATTENTION: calling new module /usr/sbin/alt_disk_copy. Please see the alt_disk_copy man page and documentation for more details. Executing command: /usr/sbin/alt_disk_copy -d "hdisk1" +-----+ --+ Calling mkszfile to create new /image.data file. Checking disk sizes. Creating cloned rootvg volume group and associated logical volumes. Creating logical volume alt_hd5. Creating logical volume alt_hd6. Creating logical volume alt_hd8. Creating logical volume alt_hd4. Creating logical volume alt_hd2. Creating logical volume alt_hd9var. Creating logical volume alt_hd3. Creating logical volume alt_hd1. Creating logical volume alt_hd10opt. Creating logical volume alt_sysdump_dev. Creating /alt_inst/ file system. Creating /alt_inst/home file system. Creating /alt_inst/opt file system. Creating /alt_inst/tmp file system. Creating /alt_inst/usr file system. Creating /alt_inst/var file system. Generating a list of files for backup and restore into the alternate file system...</pre>

Step	Action
	Backing-up the rootvg files and restoring them to the alternate file system... Modifying ODM on cloned disk. Building boot image on cloned disk. forced unmount of /alt_inst/var forced unmount of /alt_inst/usr forced unmount of /alt_inst/tmp forced unmount of /alt_inst/opt forced unmount of /alt_inst/home forced unmount of /alt_inst forced unmount of /alt_inst Changing logical volume names in volume group descriptor area. Fixing LV control blocks Fixing file system superblocks... Bootlist is set to the boot disk: hdisk1
2	Reboot the server.
3	Use the bootlist command to create a new boot list that boots only from the primary paths of the boot LUN. Example: The following example contains commands to ◆ Create a boot list in normal mode that says to boot off logical volume hd5 on disk hdisk1 and displays information about the boot. ◆ Return information on the last boot device. ◆ Display a 1 if the IPL code in the ROS on the machine running the command is capable of booting from a specified disk. <pre>-> bootlist -m normal -o hdisk1 blv=hd5 -> bootinfo -b hdisk1 -> bootinfo -B hdisk1 1</pre>
4	Use the ipl_varyon command to varyon the root volume group. <pre>5300-05-00] root@ibmbc-fak01-b4 / -> ipl_varyon -i PVNAME BOOT DEVICE PVID VOLUME GROUP ID hdisk0 YES 002006ca49c93e760000000000000000 002006ca0000d000 hdisk1 YES 002006cac6d2a0fe00000000000000000 002006ca0000d000</pre>

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

Topics covered in this chapter This chapter discusses the following topics:

- ◆ [“Troubleshooting checklists”](#) on page 98
- ◆ [“Confirming that your system is correctly set up”](#) on page 100
- ◆ [“Installation fails because system does not meet requirements”](#) on page 106
- ◆ [“Commands you can use to troubleshoot problems”](#) on page 107
- ◆ [“SAN Toolkit diagnostic utilities”](#) on page 108
- ◆ [“Change the AIX HBA’s attributes”](#) on page 114
- ◆ [“Change the value for the queue depths on the disk”](#) on page 118
- ◆ [“Change the value for the queue depths on the disk”](#) on page 118
- ◆ [“General troubleshooting”](#) on page 119

Troubleshooting checklists

Connectivity problems

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

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

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

Storage setup checklist

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

- ◆ The storage system FCP cfmodes is correctly set for your configuration. See Chapter 3, “[Migrating a Storage Cluster to Single Image Mode](#),” on page 39.
- ◆ FCP is running and all the target ports are in a good state.
- ◆ The igroups have the ostype set to AIX.
- ◆ The LUNs are
 - ❖ mapped to host igroups
 - ❖ in a good state as indicated by the output from the `sanlun lun show -p` command or the `lspath` command.
 - ❖ use AIX as their operating system type

Host setup checklist

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

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

- ◆ The host is seeing all of the LUNs and LUN paths that it should.
You can confirm this by running the `lsdev` command or, if you have multipathing, `sanlun lun show -p`.
- ◆ The paths priorities have been set. If your environment does not use ALUA, you must run the `dotpaths` utility to set the path priorities after having the host discover the LUNs.
- ◆ Use MPIO 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

Confirming that your system is correctly set up

About the sections that follow

If you have any problems with the Host Utilities, you should make sure your system setup is correct.

The information in the sections that follow helps you check your system setup. As you go through it, keep the following in mind:

- ◆ As an experienced AIX user, you may already know these commands. This is just a quick refresher of some of the key areas to check to avoid having problems.
For more information on any of the AIX or VIO server commands, see the man pages and operating system documentation.
For more information on the Data ONTAP commands, see the Data ONTAP documentation, in particular, the FCP and iSCSI Block Access Guide.
- ◆ You perform some of these checks from the host and others from the storage system. In some cases, you must have the Host Utilities SAN Toolkit installed before you perform the check. For example, the SAN Toolkit contains both the `dotpath` utility and the `san lun` command.
- ◆ To make sure you have the current version of the system components, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> or the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. Support for new components is added on an on-going basis. This online information contains a complete list of supported HBAs, platforms, applications, and drivers. It is updated on a regular basis to ensure that they always have information on the most current components.

Note

The examples that follow were accurate at the time this document was produced. They are not intended to show all the options available.

Check your operating system version

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

```
oslevel -s
```

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

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

```
# oslevel
1.4.0.0
```

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 correct host HBA is installed

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

Step	Action
1	To display HBA information, enter: <pre>lsdev -Cc adapter</pre> Example: 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 sissscsia0 Defined 04-08 PCI-X Dual Channel Ultra320 SCSI Adapter sissscsia1 Available 08-08 PCI-X Dual Channel Ultra320 SCSI Adapter usbhcc0 Available 02-08 USB Host Controller (33103500) usbhcc1 Available 02-09 USB Host Controller (33103500) vsa0 Available LPAR Virtual Serial Adapter vsa1 Available LPAR Virtual Serial Adapter</pre>
2	If you need to install an HBA, refer to the documentation that came with your HBA.

Verifying that the storage system is configured

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

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

Note

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

6. Set up ALUA, if you have Data ONTAP 7.2 or later.

Refer to “[Where to go for more information](#)” on page 10 for a list of useful documents and web sites.

Check the ALUA settings on the storage system

If you are running a version of Data ONTAP that supports ALUA, you must have ALUA set on the storage system to work with igroups. To make sure you have ALUA set for your igroup, use the `igroup show -v` command.

Example: The following command line displays information about the `cfmode` on the storage system and shows that ALUA is enabled. (To make the information on ALUA easier to locate, it is shown in bold.)

```
filerA# igroup show -v
filerA (FCP):
  OS Type: aix
  Member: 10:00:00:00:c9:4b:e3:42 (logged in on: 0c)
  Member: 10:00:00:00:c9:4b:e3:43 (logged in on: vtic)
  ALUA: Yes
```

For more information on setting up and working with ALUA, see the *Data ONTAP® Block Access Management Guide for FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml. In particular, see the section “*Enabling ALUA for a Fibre Channel igroup.*”

Setting up ALUA after installing the Host Utilities

If you need to set up ALUA after you have installed the Host Utilities, you must remove the existing NetApp LUN handles for the device. To do this, enter the following command for each device:

```
rmdev -dl device_name
```

Now you can set up ALUA.

Example: This example removes the `hdisk2` LUN handles.

```
# rmdev -dl hdisk2
```

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 the Host Utilities. See the *Fibre Channel Configuration Guide* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/FCPCConfigurationGuide.pdf.
- ◆ Turned on power to the switch.
- ◆ Configured the network parameters for the switch, including its serial parameters and IP address.

Determining whether to use switch zoning

If you have a fabric-attached configuration, determine whether switch zoning is appropriate for your system setup.

Zoning requires more configuration on the switch, but it results in the following:

- ◆ Configuring the host is simplified.
- ◆ Information becomes more manageable. The output from host tools such as `iostat` and `lsdev` is easier to read because fewer paths are displayed.

To have a high-availability configuration, make sure that each LUN has at least one primary path and one secondary path through each switch. For example, if you have two switches, you would have a minimum of four paths per LUN.

It is recommended that your configuration have no more than eight paths per LUN. For more information about zoning, see the *Configuration Guide for NetApp FCP and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Correct order for powering up the equipment

The different pieces of hardware communicate with each other in a prescribed order, which means that problems occur if you turn on power to the equipment in the wrong order. Throughout the procedures that follow, this is the order you must use when powering on the equipment:

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

Verifying the host and storage system can communicate

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

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

Installation fails because system does not meet requirements

Installation checks for system requirements

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

Reinstalling the software

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

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

Commands you can use to troubleshoot problems

Tools you can use

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

- ◆ The SAN Toolkit `sanlun` utility
- ◆ The AIX `lsdev` command

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 60.

Using the `lsdev` command

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

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

SAN Toolkit diagnostic utilities

Diagnostic utilities

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

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

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

- ◆ `filer_info`. See “[Collecting storage system information](#)” on page 109 for more information.
- ◆ `aix_info`. See “[Using the AIX data-collection tool](#)” on page 110.
- ◆ `ontap_collect`. This is a snap plug-in that you use in a VIO environment instead of `aix_info`. See “[Using the AIX data-collection tool](#)” on page 110
- ◆ `<switch>_info`, where `<switch>` is the brand of switch you are using (for example, `brocade_info`). See “[Collecting switch information](#)” on page 111.

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.

Note

In VIO server and client environments, you run the `ontap_collect` tool instead of `aix_info`. The `ontap_collect` tool is a snap plug-in that creates a file you can send to Customer Support. You execute it the same way you execute any other snap plug-in. For example, you can run it by entering a `snap ontap_collect` command.

To execute `aix_info`, complete the following steps:

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

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

Collecting switch information

If you are having a problem, Customer Support might ask you to run the switch utility the SAN Toolkit supplies when you install the toolkit software. The SAN Toolkit provides a utility for each switch it supports for this Host Utilities. 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 <code>/opt/NetApp/santools/bin</code> : <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 value for the queue depths on the disk](#)” on page 118).

Changing the FCP HBA’s attributes and rebooting

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

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

Changing the FCP HBA's attributes without rebooting

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

Note

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

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

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

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

Change the value for the queue depths on the disk

Setting the value of the queue depths

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

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

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

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

Note

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

Example of changing the queue depth

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

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

General troubleshooting

Basic operations take a long time

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

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

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

Index

A

AIX host

- example of successful installation 25
- example of successful uninstallation 37
- uninstalling Host Utilities software 32
- using SMIT 19

ALUA

- Asymmetric Logical Unit Access 5
- checking on igroup 103
- required for storage system 5

Attach Kits

- previous kits 2

attributes

- changing and rebooting 114
- changing on HBA 114
- changing with SMIT 114
- changing without rebooting 115

C

CD-ROM

- getting software 15

cfgmgr command 60, 62, 66

cfmode

- correct setting 40

chdev command 114

checklists

- filer setup 98
- host setup 98
- installation tasks 7
- troubleshooting 98

D

diagnostic utilities

- aix_info 108
- executing 108
- filer_info 108, 109
- switch_info 108, 111

disk

- changing value of queue depths 118

disk shelves

- powering up 104

dotpaths 3

- executing on VIO servers 66
- set path priorities 64, 70

E

equipment

- powering up 104

examples

- checking ALUA on igroup 103

F

FCP AIX Host Utilities

- completing installation 54
- contents 2
- defined 2
- SAN Toolkit 3

file system

- creating 75

filer_info

- diagnostic script 3
- executing 109

files

- filer_info 109
- host settings file 15
- SAN Toolkit file 15
- switch_info 111

filesset

- check version 101

H

HBA

- changing attributes 114
- installing 102
- verifying queue depth 31

host

- check OS version 100
- communicating with storage system 105
- creating volume group 73

- device for physical volume 65, 72
- discovering new LUNs 60, 62, 66
- displaying LUN information 79
- powering up 104
- setup checklist 98

- host settings file 15

- Host Utilities

- dotpath utility 64, 70
 - software packages 15
 - VIO server 4

- hosts

- VIO server 4

I

- igroup

- checking ALUA setting 103
 - definition 54

- install command

- checks system requirements 106
 - rerunning 106

L

- lsdev command 115

- using for troubleshooting 107

- lslpp command 101

- LUN ID

- defined vii

- LUNs

- best practices for creating 55
 - defined vii
 - discovering 60, 62, 66
 - displaying host information 79
 - tasks for creating on filer 55

M

- MPIO

- using on host 2

- MPxIO

- transparent paths 5
 - using with sanlun 5

- multipathing

- MPxIO 5

N

- NOW site

- downloading software 16

O

- oem_setup_env command
- using 69

- ontap_collect

- configuration script 3

- oslevel command 100

P

- padmin

- VIO servers 4

- path priorities

- setting 64, 70

- paths

- not visible 5

- physical volume

- selecting device 65, 72

- powering up equipment 104

Q

- queue depth

- changing value on disk 118

- verifying 31

S

- SAN boot

- pre-requisites 87

- SAN Toolkit 3

- components 3

- san_version command 3

- sanlun utility 3

- software package 15

- san_version

- command 3

- sanlun

- part of SAN Toolkit 3

- paths not visible 5

- sanlun fcp show adapter command 82
- sanlun lun show command 79
- sanlun utility
 - about 58
 - sanlun fcp show adapter 82
 - using for troubleshooting 107
- scripts
 - filer_info 3
 - ontap_collect 3
- set path priorities 3
- SMIT
 - creating file system 75
 - creating volume group 73
 - unloading child devices 114
 - using on AIX 19
- software
 - downloading 16
 - getting 15
 - getting from CD-ROM 15
 - installation tasks 18
 - uncompressing 17
- software packages 15
- storage
 - accessing on volume group 75
- storage system
 - communicating with host 105
 - correct dfmode 40
 - powering up 104
 - setup checklist 98
 - verifying setup 102
- switch_info
 - executing 111

- switches
 - powering up 104
 - verifying setup 104
 - zoning 104

T

- Terminology vii
- troubleshooting
 - checklists 98
 - using lsdev command 107
 - using sanlun 107

U

- uninstalling
 - Host Utilities software 32

V

- VIO servers
 - oem_setup_env command 69
 - padmin 4
 - using on host 2
 - using with Host Utilities 4
- volume group
 - accessing storage 75
 - creating 73

Z

- zoning 104

