

FC AIX® Host Utilities 5.0 for Native OS and Veritas® Installation and Setup Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
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 215-04287_A0
December 2008

Copyright and trademark information

Copyright information

Copyright © 1994-2009 NetApp, Inc. All rights reserved. Printed in the U.S.A.

The third-party copyright statements and licenses that apply to open source software components distributed with the product are contained in the section “Third-Party Copyrights and Licenses.”

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

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

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, VFM (Virtual File Manager), and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by

Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; Virtual File Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. and/or other countries.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Installation and Setup Overview	1
	Host Utilities overview	2
	Supported AIX, Data ONTAP features, configurations	5
	The stacks supported by the Host Utilities	8
	FC AIX Host Utilities installation overview	12
	Where to go for more information	17
Chapter 2	Installing the Host Utilities Software.	21
	Completing the pre-installation tasks.	22
	Configuring the target HBA for direct-attached configuration.	23
	Obtaining the software packages	24
	Installing the Host Utilities software packages.	27
	Uninstalling the Host Utilities software packages	40
Chapter 3	Veritas: Installing the ASL and APM	47
	Overview of the Array Support Library and Array Policy Module	48
	Installing the ASL and APM	50
	Uninstalling the ASL and the APM	55
	Example of uninstalling and installing the ASL and APM.	56
	Obtaining multipathing information	60
	Optimizing the Veritas DMP Restore Daemon.	65
	Understanding ASL error message types.	67
Chapter 4	Basic Tasks for Working with LUNs.	69
	Overview of LUN configuration and management.	70
	Creating and mapping igroups and LUNs	72

	High level methods for discovering LUNs	74
Chapter 5	Veritas: Configuring and Managing LUNs	77
	Discovering LUNs on a host using VxVM and DMP	78
	Working with LUNs and VxVM volumes	81
Chapter 6	Native OS, PowerVM: Configuring LUNs.	85
	Discovering LUNs with native AIX LVM (no multipathing)	86
	(Native OS only) Discovering LUNs with MPIO software	88
	(PowerVM only) Discovering LUNs with MPIO on VIO servers	92
	(Native OS only) Creating volume groups and file systems	98
Chapter 7	Using sanlun to Display LUN, HBA Information	103
	Using sanlun to display LUN information	104
	Using sanlun to display host HBA information	109
Chapter 8	Troubleshooting	111
	Troubleshooting checklists	112
	Confirming that your configuration is set up correctly.	114
	Installation fails because system does not meet requirements	120
	Tools you can use to troubleshoot problems	121
	SAN Toolkit diagnostic utilities	123
	Changing HBA attributes such as queue depth.	130
	Changing the HBA queue depth with no LUNs mapped.	135
	Changing the value for the queue depths on the disk.	136
	General troubleshooting	137
Appendix A	(Native OS) Configuring SAN booting.	139
	SAN boot overview.	140
	Setting up a SAN Boot LUN	142

Cloning the root volume group 149

Appendix B

Third-Party Copyrights and Licenses 151

Index 177

Preface

About this guide

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

See the NetApp Interoperability Matrix Tool, available on the NetApp on the Web™ (NOW) site, <http://now.netapp.com/matrix/mtx/login.do>, for current information about the system requirements.

This document applies to host environments running Multiple Path I/O (MPIO) and to host environments running Veritas Storage Foundation™. In addition, it supports host environments running the PowerVM (VIOS) environment.

Audience

This guide is for system installers and administrators who are familiar with the following:

- ◆ AIX operating systems
- ◆ MPIO or Veritas environments
- ◆ VIO server environments
- ◆ Configuring and administering a storage system
- ◆ Using FC 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*.
- ◆ *FC* refers to Fibre Channel.
- ◆ *igroup* refers to a collection of unique identifiers, either FC WWPNs (World Wide Port Names) 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 FC 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 <i>fctest adaptername</i> 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.

Attention

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

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 docomments@netapp.com. To help us direct the comments to the correct division, please put FC AIX in the subject line.

About this chapter

This chapter provides a brief overview of the FC AIX Host Utilities for Native OS and Veritas.

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 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
- ◆ [“Supported AIX, Data ONTAP features, configurations”](#) on page 5
- ◆ [“The stacks supported by the Host Utilities”](#) on page 8
- ◆ [“FC AIX Host Utilities installation overview”](#) on page 12
- ◆ [“Where to go for more information”](#) on page 17

Host Utilities overview

Important information about Host Utilities

The Host Utilities provide components that enable users to connect AIX hosts to NetApp storage systems. They support several different AIX environments:

- ◆ Veritas. This environment uses Veritas Storage Foundation to provide volume management and multipathing. For multipathing, it uses Veritas Dynamic Multipathing (DMP).
- ◆ Native OS. This environment uses features of the AIX operating system to provide volume management and multipathing. For multipathing, it uses Multiple Path I/O (MPIO) with IBM's default Path Control Module (PCM).
- ◆ PowerVM. This environment uses the Virtual I/O (VIO) Server/Client. You must be running the Native OS stack to use PowerVM; you cannot use PowerVM and Veritas together.

Note

The term “VIO” was used to refer to PowerVM in earlier documentation.

To determine which environments the Host Utilities support, check the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Note

Previous versions of the Host Utilities were called FC AIX Attach Kits.

The tools contained in the Host Utilities apply to numerous components. Sometimes the same tools are used by Host Utilities for other operating systems. As a result, the Host Utilities version number and the tools version number do not match, even though the tools are part of that Host Utilities release. This is not issue.

Overview of the FC AIX Host Utilities

The Host Utilities consist of the following components:

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

Note

You are not limited to using the Host Utilities tools to configure the host. You can also use HBA-provided tools or even a combination of the Host Utilities tools and tools provided by HBA vendors.

- ❖ The `san lun` utility that displays information about LUNs and the host HBA.
- ❖ Diagnostic scripts that provide diagnostic information about components in your configuration. Customer Support might ask for output from these scripts.
- ◆ Documentation that describes how to configure the AIX host. The documentation includes this setup guide, *Release Notes*, a quick command reference, and a guide to the host settings changed by the Host Utilities.

NetApp qualifies new components on an ongoing basis. For the most current information on the supported AIX Host Utilities stacks, see “[Where to go for more information](#)” on page 17.

Software tools provided by the Host Utilities

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

- ◆ `aix_info` script. This script collects configuration information about your host.

Note

You use this script on AIX hosts that are running Native OS or Veritas environments. Do not use this script in PowerVM environments. Not all FC AIX Host Utilities stacks support PowerVM environments.

- ◆ (Native OS only) `dotpaths` utility. This utility lets you set path priorities in environments that are not using Asymmetric Logical Unit Access (ALUA).
- ◆ `controller_info` script and `filer_info`. This script collects information about the storage system.

Note

At this time the Host Utilities support both the `controller_info` script and the `filer_info` script. Future releases of the Host Utilities may only support the `controller_info` script.

- ◆ (PowerVM environment only) `ontap_collect`. This is a snap extension that collects configuration information about your host just as the `aix_info` script does. VIO servers require you to run as the prime administrator (`padmin`) instead of `root`. To execute `ontap_collect`, you must be 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 the host. In a PowerVM environment, you can only run the `sanlun` utility on the server; you cannot run it on a VIO client.
- ◆ Diagnostic scripts. In addition to the `*_info` scripts, you can use the following scripts 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.

Remove previous versions before installing Host Utilities

This version of the FC AIX Host Utilities 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 Host Utilities software packages](#)” on page 40.

Supported AIX, Data ONTAP features, configurations

Features supported by the Host Utilities

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

- ◆ FC and iSCSI within the same AIX host and to the same target as long as they use different volume groups.
- ◆ Multiple paths to the storage system when a multipathing solution is installed
- ◆ Dynamic Logical Partitioning
- ◆ (Native stack only) Asymmetrical Logical Unit Access
- ◆ Cluster failover using single image mode only
- ◆ PowerVM Live Partition Mobility on IBM System p
- ◆ IBM General Parallel File System (GPFS)
- ◆ VIO server/client environment
- ◆ Veritas Storage Foundation
- ◆ SAN booting

For information on which features are supported with which configurations, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

About using multiple protocols on a single host

The Host Utilities support using both FC and iSCSI on the same host and pointing to the same target. Because you need to have a separate volume group for each protocol, you must use a separate LUN for each protocol. You cannot mix FC and iSCSI on the same volume group.

Note

This document discusses using Host Utilities in supported FC environments. For information on setting up the Host Utilities in an iSCSI environment, see the *iSCSI AIX Host Utilities Setup and Installation Guide*.

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. See your HBA documentation for information on installing the HBAs.

About multipathing

The FC AIX Host Utilities support different multipathing solutions based on your stack. Using multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, FC traffic continues on the remaining paths. In order for a host to have multiple paths to a LUN, you must have multipathing enabled.

The SAN Toolkit `sanlun` command will show you the path policy to which the host has access.

Dynamic LPAR is available

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

(Native OS only) ALUA support with certain versions of Data ONTAP

The FC AIX Host Utilities for Native OS support both storage systems that have ALUA enabled and those that don't.

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

To use ALUA, the storage system must be running a version of Data ONTAP® software that supports ALUA. Check the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do> to see which Data ONTAP software versions support ALUA and are supported by the Host Utilities. NetApp introduced ALUA support with Data ONTAP 7.2 and Single Image cluster failover mode (cfmode).

The supported cluster failover mode

If you upgrade to Data ONTAP 7.2 or later and have an active/active configuration (also known as a *storage cluster*), you must migrate the cfmode setting of both systems in the configuration to single image mode. For information on doing this, see *Changing the Cluster cfmode Setting in Fibre*

Channel SAN Configurations, which is available online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SICFMODE_1205.pdf.

About GPFS

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

The Host Utilities support GPFS in certain configurations. Check the *Release Notes* to see if your configuration supports this feature.

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

Note

At the time this document was prepared, IBM provided a GPFS library of documents at this location:

<http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/index.jsp?topic=/com.ibm.cluster.gpfs.doc/gpfsbooks.html>.

About Live Partition Mobility from PowerVM 1.5.2

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

The Host Utilities support PowerVM Live Partition Mobility in certain configurations. Check the *Release Notes* to see if your configuration supports this feature.

Note

At the time this document was prepared, IBM provided a PowerVM Live Partition Mobility documents at this location:

<http://www14.software.ibm.com/webapp/set2/sas/f/pm/component.html>.

About SAN booting

SAN boot is using a SAN-attached disk, like a LUN, as a boot device for an AIX host. The Host Utilities support SAN booting in Native OS environments.

The stacks supported by the Host Utilities

Finding the instructions for your AIX environment

Many of the instructions in this manual apply to all the AIX environments that the Host Utilities support. In some cases, the commands you use vary based on your AIX environment.

The following table summarizes some of the areas where you must use instructions specific to your AIX environment.

Task	Veritas stack	Native OS stack	PowerVM stack
Executing commands	Log in as root and use standard AIX and Veritas commands	Log in as root and use standard AIX commands	Run as <code>padmin</code> and use VIO commands. For some Host Utilities tasks, you must log in as <code>padmin</code> and enter the <code>oem_setup_env</code> command to change to root. See “Using a PowerVM environment with the Host Utilities (Native OS only)” on page 10.
Installing the Host Utilities host settings file for your environment	Use the file that is extracted to the SAN directory. This file provides the host setting used in Veritas environments. You must also install the SAN Toolkit software, which is in the <code>SAN_Tool_Kit</code> directory. See “Installing the Host Utilities software packages” on page 27.	Use the file that is extracted to the MPIO directory. This file provides the host setting used in Native OS and PowerVM environments. You must also install the SAN Toolkit software, which is in the <code>SAN_Tool_Kit</code> directory. See “Installing the Host Utilities software packages” on page 27.	

Task	Veritas stack	Native OS stack	PowerVM stack
Special setup for your environment	You must install the ASL. In addition, if you are using Veritas Storage Foundation 5.0, you must also install the APM. See “ Veritas: Installing the ASL and APM ” on page 47.	If you are using a version of Data ONTAP that supports ALUA, you must have ALUA running. See “ Check the ALUA settings on the storage system ” on page 117.	
Working with LUNs	See “ Veritas: Configuring and Managing LUNs ” on page 77 for information on which commands to use when working in a Veritas environment.	See “ Native OS, PowerVM: Configuring LUNs ” on page 85 for information on which commands to use when working in a Native OS environment.	See “ (PowerVM only) Discovering LUNs with MPIO on VIO servers ” on page 92 for information on which commands to use when working in a PowerVM environment.
SAN booting	Not supported.	See “ (Native OS) Configuring SAN booting ” on page 139 for more information.	

Using a Veritas environment with the Host Utilities

The Veritas environment uses the commands and tools that are provided with the Veritas Storage Foundation Suite. For this environment you must install the Veritas Storage Foundation and the Array Support Library and Array Policy Module.

You must be root both to install the Host Utilities and then execute the sanlun utility.

To manage LUNS and volumes, you use the Veritas Volume Manager (VxVM) to manage the LUNs and volumes.

For multipathing, you use Veritas Dynamic Multipathing (DMP).

Using a Native OS environment with the Host Utilities

The Native OS environment uses the commands and tools that are provided with the AIX operating system.

To install the Host Utilities and execute the sanlun utility, you run as root.

To manage the LUNs and volumes, use the Native AIX Logical Volume Manager (LVM).

For multipathing, use MPIO. If your version of Data ONTAP supports ALUA, you can use ALUA to set the paths. If not, use the `dotpaths` utility that the Host Utilities provide.

Using a PowerVM environment with the Host Utilities (Native OS only)

If you are using AIX with a VIO server, make sure you have your PowerVM environment set up before you install the Host Utilities. PowerVM 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 from the ones you use when running the Host Utilities on an AIX host using only the native stack. For example, to discover LUNs on VIO servers, you use the `cfgdev` command; on AIX hosts without PowerVM, you use the `cfgmgr` command.

Some other differences include the PowerVM:

- ◆ Installing the Host Utilities in PowerVM 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. While in this mode, you can perform activities such as installing the Host Utilities or running SAN Toolkit utility commands such as `dotpaths` or `sanlun`. To return to the normal operating mode for PowerVM environments, halt the `oem_setup_env` command mode by entering `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.
- ❖ If you have LUNs presented to a VIO server from multiple third-party storage vendors, make sure that all the LUNs use the same maximum transfer size. A good way to do this is to pick the highest value for all the LUNs and then match the others to that value using the `chdev` command.

Example: In this example, hdisk61 has a maximum transfer size of 256KB (4000h).

```
#-> lsattr -El hdisk61 -a max_transfer  
max_transfer 0x40000 Maximum TRANSFER Size True
```

- ◆ 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 applicable IBM documentation. If you want to see which commands you can use with VIO servers, enter the `help` command and then use the `man` command to see details on each command.

FC AIX Host Utilities installation overview

Overview of Installing the Host Utilities and setting up your system

Installing the FC 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. For details on these tasks,

Note

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

Prerequisites

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

1. Set up your configuration:

- ❖ Host OS, technology levels, and appropriate updates
- ❖ HBAs
- ❖ Drivers
- ❖ Volume management and multipathing, if used
- ❖ **Veritas environment only:** Veritas Storage Foundation, the Array Support Library (ASL) for the storage controllers, and, with Veritas Storage Foundation 5.0, Array Policy Module (APM)

Note

You must install the Veritas Volume Manager (VxVM) before you install the ASL and APM software. These components are available from the Symantec web site.

- ❖ Storage system and Data ONTAP
- ❖ Switches, if used

For information about the system requirements, see “[Where to go for more information](#)” on page 17.

2. Verify that your storage system is:

- ❖ Licensed correctly for the FC service and running it
- ❖ Using the recommended `cfmode` (single-image)
- ❖ Configured to work with the target HBAs

- ❖ Set up to work with the AIX host and the initiator HBAs
- ❖ Set up to work with ALUA, if it is supported with your version of Data ONTAP

Note

ALUA is not supported with Veritas Storage Foundation.

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

Note

You should 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
 - ❖ Cabled according to the instructions that came with the switch
4. Verify that:
 - ❖ The equipment is powered on in the correct order: switches, disk shelves, storage systems, and then the host.
 - ❖ The host and the storage system can communicate

Installation overview

Installing the Host Utilities involves the following tasks:

1. Get a copy of the compressed Host Utilities file, which contains the host settings package for your multipathing solution and the SAN Toolkit software package.
 - ❖ Download the compressed file containing the packages from the NOW site.
 - ❖ Copy the compressed file from the CD.
 - ❖ Extract the software packages from the file. There will be three packages, but you only need to install two of them: the host settings package for your multipathing solution (either MPIO or non-MPIO) and the SAN Toolkit package.
2. Install the software packages:
 - ❖ Login as root.

PowerVM environment only: If you are on a VIO server, login as padmin and then use the `oem_setup_env` command to change to root.
 - ❖ Go to the directory containing the extracted software packages.

- ❖ Use either the AIX System Management Interface Tool (SMIT) or the `installp` command to install the host settings software and the toolkit software.

PowerVM environment only: If you are using a PowerVM environment, you should install the SAN Toolkit software on the server and 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. Check your cluster failover mode (cfmode):
 - ❖ If you have an active/active configuration, migrate your cfmode setting to single image mode.
 - ❖ If you have a direct-attached storage, verify that the media type is correct.

Configuration overview

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. You will need to supply a World Wide Port Name (WWPN) for each of the host's HBAs.
- ◆ **Native OS, PowerVM environments only:** Enable ALUA, if you haven't already done so.
- ◆ **Native OS, PowerVM environments only:** Set 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.
 - ❖ **Native OS, Veritas environments only:** Use the `cfgmgr` command.
 - ❖ **PowerVM environment only:** Use the `cfgdev` command.
- ◆ **PowerVM environment only:** If you have LUNs presented to a VIO server from multiple third-party storage vendors, make sure that all the LUNs use the same maximum transfer size. Use the `lsatr -El <disk> -a max_transfer` 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.

For detailed information on these tasks

The chapters in this document provide more details about setting up your system, verifying your setup before you install the Host Utilities, installing and configuring the Host Utilities, creating igroups and LUNs, and displaying information on the LUNs and HBAX. The following list summarizes some of the information in the chapters. For details on what is in each chapter, go to that chapter.

- ◆ Chapter 2, “[Installing the Host Utilities Software](#),” on page 21
 - ❖ “[Completing the pre-installation tasks](#)” on page 22
 - ❖ “[Obtaining the software packages](#)” on page 24
 - ❖ “[Installing the Host Utilities software packages](#)” on page 27
 - ❖ “[Verifying the HBA initiator queue depth](#)” on page 39
- ◆ Chapter 3, “[Veritas: Installing the ASL and APM](#),” on page 47
 - ❖ “[Installing the ASL and APM](#)” on page 50
 - ❖ “[Example of uninstalling and installing the ASL and APM](#)” on page 56
 - ❖ “[Obtaining multipathing information](#)” on page 60
 - ❖ “[Optimizing the Veritas DMP Restore Daemon](#)” on page 65
- ◆ Chapter 4, “[Basic Tasks for Working with LUNs](#),” on page 69
 - ❖ “[Creating and mapping igroups and LUNs](#)” on page 72
 - ❖ “[High level methods for discovering LUNs](#)” on page 74
- ◆ Chapter 5, “[Veritas: Configuring and Managing LUNs](#),” on page 77
 - ❖ “[Discovering LUNs on a host using VxVM and DMP](#)” on page 78
 - ❖ “[Working with LUNs and VxVM volumes](#)” on page 81
- ◆ Chapter 6, “[Native OS, PowerVM: Configuring LUNs](#),” on page 85
 - ❖ “[Discovering LUNs with native AIX LVM \(no multipathing\)](#)” on page 86
 - ❖ “[\(Native OS only\) Discovering LUNs with MPIO software](#)” on page 88
 - ❖ “[\(PowerVM only\) Discovering LUNs with MPIO on VIO servers](#)” on page 92
 - ❖ “[\(Native OS only\) Creating volume groups and file systems](#)” on page 98

- ◆ Chapter 7, [“Using sanlun to Display LUN, HBA Information,”](#) on page 103
 - ❖ [“Using sanlun to display LUN information”](#) on page 104
 - ❖ [“Using sanlun to display host HBA information”](#) on page 109
- ◆ Chapter 8, [“Troubleshooting,”](#) on page 111
 - ❖ [“Confirming that your configuration is set up correctly”](#) on page 114
 - ❖ [“Tools you can use to troubleshoot problems”](#) on page 121
 - ❖ [“SAN Toolkit diagnostic utilities”](#) on page 123
 - ❖ [“Changing HBA attributes such as queue depth”](#) on page 130
 - ❖ [“Changing the HBA queue depth with no LUNs mapped”](#) on page 135
 - ❖ [“Changing the value for the queue depths on the disk”](#) on page 136
- ◆ Optional: Appendix A, [“\(Native OS\) Configuring SAN booting,”](#) on page 139

In addition, the section [“Where to go for more information”](#) on page 17 provides pointers to other documents and Web sites that can help you in working with your system.

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/ Note— The <i>Release Notes</i> are updated more frequently than the rest of the documentation. You should always check the release notes before installing the Host Utilities to see if there have been any changes to the installation or setup process since this document was prepared. You should also check them periodically to see if there is new information on using the Host Utilities. A summary of what has been updated and when is on the <i>Release Notes</i> index page.
The latest supported Host Utilities configurations and environments	◆ <i>NetApp Interoperability Matrix Tool</i> at http://now.netapp.com/matrix/mtx/login.do ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
A summary of some of the commands you might use with the Host Utilities	The <i>Quick Reference</i> for the FC 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
Changes to the host settings that are recommended by the Host Utilities	<i>Host Settings Affected by the Host Utilities</i> at http://now.netapp.com/NOW/knowledge/docs/san/
Supported Fibre Channel SAN topologies	<i>FCP Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

If you need information about...	See the following product documentation...
Changing the cfmode	<i>Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf
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 ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Enabling ALUA	◆ <i>Data ONTAP Block 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 ISCSI and 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

If you need information about...	See the following product documentation...
Changing the cfmode	<i>Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations</i> at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf
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 ◆ <i>System Configuration Guide</i> at http://now.netapp.com/NOW/knowledge/docs/docs.shtml
Enabling ALUA	◆ <i>Data ONTAP Block 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 ISCSI and 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

If you need information about...	See the following product documentation...
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 Veritas	Refer to your Veritas documentation. In particular, the following documents in particular might be helpful: ◆ Veritas Storage Foundation 5.0 ❖ Veritas Storage Foundation and High Availability Solutions 5.0 - Getting Started Guide (AIX) ❖ Veritas Storage Foundation 5.0 - Installation Guide (AIX) ❖ Veritas Storage Foundation for Oracle RAC 5.0 - Installation and Configuration Guide (AIX) ◆ Veritas Storage Foundation 4.0 ❖ Veritas Storage Solutions 4.0 Getting Started Guide for AIX ❖ Veritas Storage Foundation 4.0 Installation Guide/AIX ❖ Veritas Storage Foundation 4.0 for Oracle RAC - Installation and Configuration Guide (AIX)
Working with PowerVM environments	Refer to your operating system documentation. Additional documentation on using PowerVM environments is available on the IBM web site.

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

Topics in this chapter This chapter discusses the following topics:

- ◆ [“Completing the pre-installation tasks”](#) on page 22
- ◆ [“Obtaining the software packages”](#) on page 24
- ◆ [“Installing the Host Utilities software packages”](#) on page 27
- ◆ [“Uninstalling the Host Utilities software packages”](#) on page 40

Completing the pre-installation tasks

Before you begin

You need to have your operating system, drivers, HBAs, and multipathing set up before you install the Host Utilities. You also need to have any features specific to your AIX environment set up.

- ◆ Veritas environments: You must have Veritas Storage Foundation installed and set up before you install the Host Utilities. In addition, you should configure VxVM and, if you want to use multipathing, DMP. Depending on your version of Veritas, you may need to install the ASL and the APM. For more information, see [“Veritas: Installing the ASL and APM”](#) on page 47.
- ◆ Native OS and PowerVM environments: You must have MPIO set up. If you are using ALUA, you must have that setup.
- ◆ PowerVM environments: You must have the PowerVM environment set up.

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

- ◆ Verify that your host system meets the system requirements. See [“Where to go for more information”](#) on page 17.
- ◆ If you currently have an earlier version of the FC AIX Host Utilities or FC AIX Attach Kit installed, you must remove it. See [“Uninstalling the Host Utilities software packages”](#) on page 40.

You should also have your storage system set up.

- ◆ The storage system must be able to communicate with the host.
- ◆ Active/active configurations (clustered) must use single image cfmode. If your system is not configured for single image, see *Changing the Cluster cfmode Setting in Fibre Channel SAN Configurations*, which is available online at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/QuickRef/SSICFMODE_1205.pdf.
- ◆ Direct-attached configurations must have the media type set to loop. See [“Configuring the target HBA for direct-attached configuration”](#) on page 23

For information on verifying that your system setup, see [“Confirming that your configuration is set up correctly”](#) on page 114.

Configuring the target HBA for direct-attached configuration

Correct media type for direct-attached configuration

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

Note—
If you are using an active/active configuration (cluster), the media type setting is disabled.

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

Setting the media type to loop

To set the media type to loop, perform the following steps.

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

Obtaining the software packages

Software packages required for the Host Utilities

The Host Utilities are packaged into a single compressed file. This file contains:

- ◆ Two host settings packages. There is settings package for environments that use DMP multipathing and one for environments that use MPIO. You only need to install the package for your environment. The host settings file contains NetApp's FC device definitions, which are required for FC AIX configurations. This is an Object Data Manager (ODM) stanza LPP package.
- ◆ The SAN Toolkit package. You install the same package regardless of your configuration. 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 might ask you to run the `sanlun` utility or a diagnostic script if a problem occurs.

Note

This document uses names for these packages in its examples to make the examples easier to read. The final names for these packages may differ from those used in examples. Do not be concerned if the names of the packages you actually have are different from the ones in this document. The names used on the NOW site and Host Utilities CD are correct.

Obtaining the software packages

You can obtain the software packages in two ways:

- ◆ Download the compressed file from the NOW site at <http://now.netapp.com/NOW/cgi-bin/software>. For information about downloading the file, see “[Downloading the software from the NOW site](#)” on page 25.
- ◆ Obtain the file from the FC AIX Host Utilities CD-ROM. Information on obtaining the software package files from the CD-ROM is detailed in the section “[Installing the Host Utilities software packages](#)” on page 27.

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

Obtaining the file from the CD-ROM

To obtain the compressed file from the CD-ROM, complete the following steps.

Step	Action
1	Insert the CD-ROM into the CD-ROM drive.
2	Change to the root directory of the CD-ROM: <pre># cd <CDROM_mountpoint></pre>
3	Copy the compressed file from the CD-ROM's root directory to a temporary location on the host. Example: If you mount the CD-ROM on /mnt and want to place the SAN Toolkit file in the /tmp directory, you might enter commands similar to the following: <pre># mount -v cdrfs -o ro /dev/cd0 /mnt # cp /mnt/ ntap_aix_host_utilities_5.0.tar.Z /tmp # cd /tmp</pre>

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 and select the Download Software page.
2	On the Download Software page, go to the FC Host Utilities (Attach Kits) product row of the table and select AIX from the Select Platform drop-down list and click Go.
3	Click Go.
4	Follow the prompts to reach the Software Download page.
5	Download the compressed file to a local directory on your host. Note _____ If you download the file to a machine other than the AIX host, you must move it to the host. _____

Uncompressing the software packages

To uncompress the software packages, complete the following steps.

Step	Action
1	Change to the directory containing the downloaded file and enter the following commands to uncompress it. Note _____ The name of the file you download or get from the CD-ROM may be different from the filename shown in these examples. _____ <pre># zcat ntap_aix_host_utilities_5.0.tar.Z tar -xvf -</pre> Result: The extracted software is placed in three directories: ◆ MPIO. This directory contains the host settings software package you must install when you're using a Native OS or PowerVM environment running MPIO. ◆ SAN. This directory contains the host settings host settings software package you must install when you're using a Veritas environment running DMP. ◆ SAN_Tool_Kit. This directory contains the SAN Toolkit file. You always install this file regardless of your host environment. Example: Executing the <code>zcat</code> and <code>tar</code> commands produces output similar to the following: <pre>#-> zcat ntap_aix_host_utilities_5.0.tar.Z tar xvf - x ntap_aix_host_utilities_5.0 x ntap_aix_host_utilities_5.0/SAN_Tool_Kit x ntap_aix_host_utilities_5.0/SAN_Tool_Kit/NetApp.SAN_toolkit, 1343488 bytes, 2624 tape blocks x ntap_aix_host_utilities_5.0/SAN x ntap_aix_host_utilities_5.0/SAN/NetApp.fcp_attach_kit, 24576 bytes, 48 tape blocks x ntap_aix_host_utilities_5.0/SAN/NetApp.iscsi_support_kit, 20480 bytes, 40 tape blocks x ntap_aix_host_utilities_5.0/MPIO x ntap_aix_host_utilities_5.0/MPIO/NetApp.mpio_attach_kit, 184320 bytes, 360 tape blocks</pre>
2	Go to “Installing the Host Utilities software packages” on page 27.

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.

If you have an FC 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 Host Utilities software packages](#)” on page 40.

Veritas environments: If you have SANpath installed, you must remove it. You cannot use both Veritas and SANpath.

The correct software for your environment

You must install the correct software packages for your host environment.

- ◆ For Native OS and PowerVM environments running MPIO, use the software package that is extracted to the MPIO directory.
- ◆ For Veritas environments, use the software package that is extracted to the SAN environment.

You must also install the SAN Toolkit. This software package is extracted to the `SAN_Tool_Kit` directory. You install the same software package regardless of whether you are using a Native OS, PowerVM, or Veritas environment.

Using SMIT to install the software

To install the Host Utilities, you must install both the host setting software package and the SAN Toolkit software package. To do this using SMIT, 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 PowerVM environment, you must install the SAN Toolkit on each client.

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

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

Step	Action
5	At the Install Software screen, specify the location of the software in one of the following ways: ◆ Manually enter the location by providing the following information. ❖ If you're installing from the CD-ROM, enter the CD-ROM drive. ❖ If you're installing from the host machine, enter the path to the software package (for example, if you use MPIO for multipathing, you would enter <code>/var/tmp/ntap_aix_mpio_5.0/NetApp.mpio_attach_kit</code>). ◆ 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 the name for the package you are installing, either <code>NetApp.mpio_attach_kit</code> or <code>NetApp.veritas_attach_kit</code>. 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 MPIO software package when you're at the Install Software screen. If you were using the Veritas environment, you would be installing the SAN package (<code>NetApp.SAN_attach_kit</code>). <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_5.0/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_5.0/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 FC and iSCSI. This is not a problem. If you are using only FC, ignore the iSCSI fileset. _____ Example of a successful installation: The following is an example of output you might see when an installation using the MPIO fileset successfully completes. <pre> COMMAND STATUS Command: running stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.mpio_attach_kit.pcmadm 5.0.0.0 NetApp.mpio_attach_kit.iscsi 5.0.0.0 NetApp.mpio_attach_kit.fcp 5.0.0.0 NetApp.mpio_attach_kit.config 5.0.0.0 << Copyright notice for NetApp.mpio__attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2008. 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.pcm 5.0.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.isc 5.0.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.fcp 5.0.0.0 USR APPLY SUCCESS NetApp.mpio_attach_kit.conf 5.0.0.0 USR APPLY SUCCESS COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. NetApp.mpio_attach_kit.config 5.0.0.0 << Copyright notice for NetApp.mpio_attach_kit >> (C) Copyright Network Appliance, Inc. 2003-2008. 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>

uncompress the Host Utilities file. From the directory containing the uncompressed files, go to
/ntap_aix_host_utilities_5.0/{MPIO | SAN} directory.

Note.

If you are running in a PowerVM environment, use the `oem_setup_env` command to become root.

To install the host settings software for a Native OS environment, enter the command line:

```
# installp -aXd NetApp.mpio_attach_kit NetApp.mpio_attach_kit
```

To install the host settings software for a Veritas environment, enter the command line:

```
# installp -aXd NetApp.SAN_attach_kit NetApp.SAN_attach_kit
```

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

```
# installp -aXd NetApp.SAN_toolkit NetApp.SAN_toolkit
```

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...

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


Selected Filesets
-----
NetApp.SAN_toolkit.sanlun 5.0.0.0          # Network Appliance SAN Toolki...
NetApp.SAN_toolkit.scripts 5.0.0.0        # 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 5.0.0.0
           NetApp.SAN_toolkit.sanlun 5.0.0.0
. . . . . << Copyright notice for NetApp.SAN_toolkit >> . . . . .
(C) Copyright Network Appliance, Inc. 2003-2008.
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	5.0.0.0	USR	APPLY	SUCCESS
NetApp.SAN_toolkit.sanlun	5.0.0.0	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 in a PowerVM environment.

```
vioserver$ oem_setup_env
# PS1="$(hostname)# "
vioserver# zcat ntap_aix_mpio_5.0.tar.Z | tar xvf -
x ./ntap_aix_mpio_5.0/NetApp.mpio_attach_kit, 195584 bytes, 382 media blocks.
vioserver# cd ntap_aix_mpio_5.0
vioserver# installp -aXd NetApp.mpio_attach_kit NetApp.mpio_attach_kit
```

```

+-----+
                Pre-installation Verification...
+-----+
Verifying selections...done
Verifying requisites...done
Results...

SUCSESSES
-----
    Filesets listed in this section passed pre-installation verification
    and will be installed.
Selected Filesets
-----
NetApp.mpio_attach_kit.config 5.0.0.0      # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.fcp 5.0.0.0        # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.iscsi 5.0.0.0      # Network Appliance MPIO PCM K...
NetApp.mpio_attach_kit.pcmodm 5.0.0.0     # 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.pcmodm 5.0.0.0
    NetApp.mpio_attach_kit.iscsi 5.0.0.0
    NetApp.mpio_attach_kit.fcp 5.0.0.0
    NetApp.mpio_attach_kit.config 5.0.0.0

. . . . . << Copyright notice for NetApp.mpio_attach_kit >> . . . . .
(C) Copyright Network Appliance, Inc. 2003-2008.
All rights reserved.

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

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

```

```
installp:  bosboot process completed.
```

Summaries:

Installation Summary

Name	Level	Part	Event	Result
NetApp.mpio_attach_kit.pcmo	5.0.0.0	USR	APPLY	SUCCESS
NetApp.mpio_attach_kit.iscs	5.0.0.0	USR	APPLY	SUCCESS
NetApp.mpio_attach_kit.fcp	5.0.0.0	USR	APPLY	SUCCESS
NetApp.mpio_attach_kit.conf	5.0.0.0	USR	APPLY	SUCCESS

Verifying the HBA initiator queue depth

By default, the Host Utilities software sets the value of `num_cmd_elems` to 200 for the HBAs and the value of `queue_depth` to 12 on the hdisk handles. Testing at NetApp showed that these defaults are a good starting point for the queue depths.

Verify the value of the HBA initiator `num_cmd_elems.` and queue depth. To check these values, use the `lsattr -El` command.

Example: For HBA fcs0, you would enter the following command to check of num cmd elems.

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

Uninstalling the Host Utilities 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.
- ◆ Stop all I/O on the system and export the volume groups.
- ◆ Remove the existing NetApp LUN handles.

These procedures disconnect NetApp LUNs from the host before you uninstall the software.

Methods for uninstalling the software

You can uninstall the Host Utilities software using either SMIT or the `installp` command.

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: # <code>smit</code>

Step	Action
3	At the screen that appears, select the “Software Installation and Maintenance” menu option. Example: When you start SMIT, it displays the following screen. On this screen, the “Software Installation and Maintenance” option is the first menu option. <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 (either NetApp.mpio_attach_kit or NetApp.veritas_attach_kit). (Make sure that the “Preview only” option is set to no.) ◆ Press F4 to display a list of names. If you use F4, complete the following steps: 1. Press F4. 2. Scroll down the list of names until you reach the name for the host Utilities package. 3. Select that name 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 when you are removing the Host Utilities software package for environments using MPIO for multipathing. Note By default PREVIEW only? is set to yes. You must change it to no if you want to uninstall. +-----+ 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 +-----+

Step	Action
7	You must reboot your system for the changes to take effect. Note_____ If you do not reboot your system, the changes will not take effect. _____

Using a command line to uninstall the software

You can use a command line to uninstall this software if you prefer. You must be logged on as root to do this. To uninstall the host settings software from

- ◆ Native OS and PowerVM environments running MPIO, enter the command line:

```
# installp -u NetApp.mpio*
```

- ◆ Veritas environments, enter the command line:

```
# installp -u NetApp.san*
```

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

```
# installp -u NetApp.SAN_toolkit*
```

About this chapter

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

Note

With the Veritas Storage Foundation 4.0 with MP3, the APM is distributed as a point patch to Volume Manager 4.0.3.3 (Point-Patch_3), so you do not have to perform a separate installation.

With Veritas Storage Foundation 5.0 or 4.0 with MP4, you must install both the ASL and the APM.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Overview of the Array Support Library and Array Policy Module”](#) on page 48
- ◆ [“Installing the ASL and APM”](#) on page 50
- ◆ [“Uninstalling the ASL and the APM”](#) on page 55
- ◆ [“Obtaining multipathing information”](#) on page 60
- ◆ [“Optimizing the Veritas DMP Restore Daemon”](#) on page 65
- ◆ [“Understanding ASL error message types”](#) on page 67

Overview of the Array Support Library and Array Policy Module

What the Array Support Library is

The Array Support Library (ASL) is a NetApp-qualified library that provides information about storage array attribute configuration to the Device Discovery Layer (DDL).

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

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

Note

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

Using storage with the ASL

You cannot use storage systems simultaneously as Just a Bunch of Disks (JBODs) and vendor arrays. If you install the ASL, storage systems cannot be configured in VxVM as JBODs. They are reported as storage arrays unless you explicitly exclude them by using the `vxddladm excludearray` command.

By using the ASL, you obtain the following information about the LUNs:

- ◆ Enclosure name. VxVM's enclosure-based naming feature creates disk names based on the name of its enclosure, or disk array, and not a raw device name.
- ◆ Array Type. The NetApp storage systems are accessed as an A/A-NETAPP disk array. All primary (active) paths can be used for I/O concurrently. When all primary paths to a LUN are unavailable in a clustered storage configuration, the secondary (standby) paths are used.

What the ASL provides

The ASL provides enclosure-based naming information about SAN-attached storage systems. With enclosure-based naming, the name of the disk is based on the logical name of its enclosure, or disk array, and not a raw device name. The ASL provides specific information to VxVM about SAN-attached storage systems, instead of referring to them as JBOD devices or raw devices.

What the Array Policy Module is

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

Installing the ASL and APM

Installation overview

Installing the ASL and APM involves the following tasks:

- ◆ Verifying that your configuration meets system requirements. To locate information on system requirements, see “[Where to go for more information](#)” on page 17.
- ◆ Determining the ASL version if it is currently installed.
- ◆ Obtaining the ASL and the APM.
- ◆ Installing the ASL and APM.

Determining the ASL version

If you currently have the ASL installed, use the Veritas `vxddladm listversion` command to determine its version.

Example: The following shows output from the `vxddladm listversion` command.

```
# vxddladm listversion
LIB_NAME                               ASL_VERSION           Min. VXVM version
=====
libvxshark.so                          vm-5.0-rev-1          5.0
libvxcscovrts.so                       vm-5.0-rev-1          5.0
libvxhitachi.so                        vm-5.0-rev-1          5.0
libvxibmds6k.so                        vm-5.0-rev-1          5.0
libvxvpath.so                          vm-5.0-rev-1          5.0
libvxvp256.so                          vm-5.0-rev-1          5.0
libvxibmds4k.so                        vm-5.0-rev-2          5.0
libvxemc.so                            vm-5.0-rev-2          5.0
libvxCLARiiON.so                       vm-5.0-rev-1          5.0
libvxhds.so                            vm-5.0-rev-1          5.0
libvxhdsalua.so                        vm-5.0-rev-1          5.0
libvxhds9980.so                        vm-5.0-rev-1          5.0
libvxibmds8k.so                        vm-5.0-rev-1          5.0
libvxpurple.so                         vm-5.0-rev-1          5.0
libvxhdsusp.so                         vm-5.0-rev-2          5.0
libvxhpalua.so                         vm-5.0-rev-1          5.0
libvxibmsvc.so                         vm-5.0-rev-1          5.0
libvxvscsi.so                          vm-4.0-rev-1          5.0
libvxnetapp.so                         vm-5.0-rev-1          5.0
#
```

Obtaining the ASL and APM

The ASL and APM are not distributed with the Host Utilities. You obtain the ASL and APM from the Symantec web site.

- ◆ Veritas Storage Foundation 4.0 with MP3: The APM is distributed as a point patch to Volume Manager 4.0.3.3 (Point-Patch_3) so you only need to install the ASL software once you have VxVM installed.
- ◆ Veritas Storage Foundation 5.0 and Veritas Storage Foundation 4.0 with MP4: the ASL download file contains the software packages for both the ASL and the APM. You must extract these software packages and then install each one separately.

For information on getting these software packages, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Note

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

Upgrading the ASL and APM when running Storage Foundation 5.0

If you are using Veritas Storage Foundation 5.0, you need to install the ASL and the APM. If you already have these installed from an earlier version, you must first remove them. For more information on uninstalling the ASL, see “[Uninstalling the ASL and the APM](#)” on page 55.

Installing the ASL and APM

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

Step	Action
1	Make sure you are logged into the host as the root user.
2	If you have VxVM installed, stop all applications using VxVM volumes. Make sure that there are no VxVM volumes open, mounted, or in use. Note_____ Check the <i>NetApp Interoperability Matrix Tool</i> to make sure you have the supported maintenance packages. You should install the correct VxVM and maintenance packages before you install the ASL and APM. _____

Step	Action
3	For each host HBA (<i>fscsiX</i>), enable the AIX FastFail feature. Note_____ Enabling the FastFail feature is a requirement when you Veritas Storage Foundation. The FastFail feature improves I/O failover handling in the NetApp environment. _____ Example: <pre>chdev -l fscsi0 -a fc_err_recov=fast_fail -P chdev -l fscsi1 -a fc_err_recov=fast_fail -P</pre>
4	For each host HBA (<i>fscsiX</i>), enable Dynamic Tracking. Example: <pre>chdev -l fscsi0 -a dyntrk=yes -P chdev -l fscsi1 -a dyntrk=yes -P</pre>
5	Install the ASL according to the directions provided in the ASL TechNote on the Symantec Web site. Note_____ if you have Veritas Storage Foundation 5.0 or 4.0 with MP4, you must also install the APM. _____
6	If you have Veritas Storage Foundation 4.0 with MP3, install the Volume Manager point patch that contains the APM. Use the procedures documented in the README file that ships with the point patch. If you have Veritas Storage Foundation 5.0 or 4.0 with MP4, skip this step.
7	Reboot the host. Note_____ If you do not reboot the host, the changes will not take effect. _____

Step	Action
8	Verify that FastFail and Dynamic Tracking are enabled by entering the following command for each HBA: <pre>lsattr -El fscsiX</pre> <i>X</i> is the controller/HBA ID. Output Example: The following output indicates that Dynamic Tracking is on (dyntrk value=yes) and the FastFail feature is enabled (fc_err_recov=fast_fail): <pre>attach switch How this adapter is CONNECTED False dyntrk yes Dynamic Tracking of FC Devices True fc_err_recov fast_fail FC Fabric Event Error RECOVERY Policy True scsi_id 0x10400 Adapter SCSI ID False sw_fc_class 3 FC Class for Fabric True</pre>
9	Verify that the ASL is installed by entering the following command. <pre>vxddladm listsupport</pre> Result: The libvxnetapp.so library is listed in the command output. See the example below: <pre># vxddladm listsupport LIBNAME VID ===== libvxshark.so IBM libvxcscovrts.so CSCOVRTS libvxhitachi.so HITACHI libvxibmds6k.so IBM libvxvpath.so IBM libvxxp256.so HP libvxibmds4k.so IBM libvxemc.so EMC libvxhdsusp.so HITACHI libvxhds.so HITACHI libvxhdsalua.so HITACHI libvxhds9980.so HITACHI libvxibmds8k.so IBM libvxpurple.so SUN libvxCLARiion.so DGC libvxhpalua.so HP, COMPAQ libvxibmsvc.so IBM libvxvscsi.so AIX libvxnetapp.so NETAPP</pre>

Step	Action																																																																	
10	Enter the <code>vxddmpadm listapm all</code> command to verify that the APM is installed. Result: The output displays the Array Type as A/A-NETAPP. If the storage controller disks have not yet been detected by the host, the state of the APM is Not-Active. See the output example below. <table><tr><th>Module Name</th><th>APM Name</th><th>APM Version</th><th>Array Types</th><th>State</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>dmpaa</td><td>dmpaa</td><td>1</td><td>A/A</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P</td><td>Active</td></tr><tr><td>dmpap</td><td>dmpap</td><td>1</td><td>A/P-C</td><td>Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-VERITAS</td><td>Not-Active</td></tr><tr><td>dmpapf</td><td>dmpapf</td><td>1</td><td>A/PF-T3PLUS</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG</td><td>Not-Active</td></tr><tr><td>dmpapg</td><td>dmpapg</td><td>1</td><td>A/PG-C</td><td>Not-Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>Disk</td><td>Active</td></tr><tr><td>dmpjbod</td><td>dmpjbod</td><td>1</td><td>APdisk</td><td>Active</td></tr><tr><td>dmpnetapp</td><td>dmpnetapp</td><td>1</td><td>A/A-NETAPP</td><td>Not-Active</td></tr><tr><td>dmpsvc</td><td>dmpsvc</td><td>1</td><td>A/A-IBMSVC</td><td>Not-Active</td></tr></table>	Module Name	APM Name	APM Version	Array Types	State	=====	=====	=====	=====	=====	dmpaa	dmpaa	1	A/A	Active	dmpap	dmpap	1	A/P	Active	dmpap	dmpap	1	A/P-C	Active	dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active	dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active	dmpapg	dmpapg	1	A/PG	Not-Active	dmpapg	dmpapg	1	A/PG-C	Not-Active	dmpjbod	dmpjbod	1	Disk	Active	dmpjbod	dmpjbod	1	APdisk	Active	dmpnetapp	dmpnetapp	1	A/A-NETAPP	Not-Active	dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active
Module Name	APM Name	APM Version	Array Types	State																																																														
=====	=====	=====	=====	=====																																																														
dmpaa	dmpaa	1	A/A	Active																																																														
dmpap	dmpap	1	A/P	Active																																																														
dmpap	dmpap	1	A/P-C	Active																																																														
dmpapf	dmpapf	1	A/PF-VERITAS	Not-Active																																																														
dmpapf	dmpapf	1	A/PF-T3PLUS	Not-Active																																																														
dmpapg	dmpapg	1	A/PG	Not-Active																																																														
dmpapg	dmpapg	1	A/PG-C	Not-Active																																																														
dmpjbod	dmpjbod	1	Disk	Active																																																														
dmpjbod	dmpjbod	1	APdisk	Active																																																														
dmpnetapp	dmpnetapp	1	A/A-NETAPP	Not-Active																																																														
dmpsvc	dmpsvc	1	A/A-IBMSVC	Not-Active																																																														

Next steps

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

- ◆ If you have Data ONTAP 7.2 or later, you should change the `cfmode` setting of your active/active configuration to single image and then reconfigure your host to discover the new paths to the disk.
- ◆ On the storage system, create LUNs and map them to igroups containing the WWPNs of the host HBAs. See [“Creating and mapping igroups and LUNs”](#) on page 72.
- ◆ On the host, discover the new LUNs and configure them to be managed by VxVM. See [“Discovering LUNs on a host using VxVM and DMP”](#) on page 78.

Uninstalling the ASL and the APM

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

Only one version of an ASL that supports Data ONTAP storage system can be installed on the host at any given time.

Removing the ASL and APM To remove the ASL and the APM, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Stop all I/O to the NetApp LUNs. Note _____ In a Veritas Storage Foundation RAC cluster, you must also stop clustering on a node before you remove the ASL. _____
3	Use either SMIT or the command line to remove the ASL package. For more information on removing the ASL, see the appropriate Veritas documentation.
4	Use either SMIT or the command line to remove the APM package. For more information on removing the APM, see the appropriate Veritas documentation. Note _____ For Veritas Storage Foundation 4.0, please contact Symantec for instructions on removing the APM. _____

Selected Filesets

```
-----
VRTSNTAPapm 5.0.0.0          # Veritas NetApp Array Policy ...
VRTSNTAPasl 5.0.3.0         # VERITAS NetApp Array Support...
<< End of Success Section >>
```

FILESET STATISTICS

```
-----
  2  Selected to be deinstalled, of which:
    2  Passed pre-deinstall verification
-----
```

```
  2  Total to be deinstalled
```

```
+-----+
                        Deinstalling Software...
+-----+
```

```
installp: DEINSTALLING software for:
          VRTSNTAPapm 5.0.0.0
```

```
Check if Module is loaded
Unloading the Module
Module is busy: Cannot remove the package
```

```
WARNING: deinstal: Problem executing ./VRTSNTAPapm.unconfig_d.
Continuing...
```

```
Filesets processed:  1 of 2  (Total time:  2 secs).
```

```
installp: DEINSTALLING software for:
          VRTSNTAPasl 5.0.3.0
```

```
Unloading the library
Finished processing all filesets.  (Total time:  3 secs).
```

```
+-----+
                        Summaries:
+-----+
```

Installation Summary

Name	Level	Part	Event	Result
VRTSNTAPapm	5.0.0.0	ROOT	DEINSTALL	SUCCESS
VRTSNTAPapm	5.0.0.0	USR	DEINSTALL	SUCCESS
VRTSNTAPasl	5.0.3.0	ROOT	DEINSTALL	SUCCESS
VRTSNTAPasl	5.0.3.0	USR	DEINSTALL	SUCCESS

Installing the ASL and the APM

The following is a sample installation of the ASL and the APM when you have Veritas Storage Foundation 5.0. If you were actually doing this, your output would vary slightly based on your system setup.

```
#-> installp -aXYd . VRTSNTAPapm VRTSNTAPasl
-----+
Pre-installation Verification...
-----+
Verifying selections...done
Verifying requisites...done
Results...

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

Selected Filesets
-----
VRTSNTAPapm 5.0.0.0          # Veritas NetApp Array Policy ...
VRTSNTAPasl 5.0.3.0         # VERITAS NetApp Array Support...

<< End of Success Section >>

-----+
BUILDDATE Verification ...
-----+
Verifying build dates...done
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:
      VRTSNTAPasl 5.0.3.0

. . . . . << Copyright notice for VRTSNTAPasl >> . . . . .
Copyright ? 1990-2006 Symantec Corporation. All rights reserved.
```

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

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

. << End of copyright notice for VRTSNTAPas1 >>.

```
mv libvxnetapp.so.2 libvxnetapp.so.3
ln -s /etc/vx/lib/discovery.d/libvxnetapp.so.3 /etc/vx/lib/discovery.d/libvxnetapp.so
mv libvxnetapp.key.2 libvxnetapp.key.3
ln -s /etc/vx/aslkey.d/libvxnetapp.key.3 /etc/vx/aslkey.d/libvxnetapp.key
Adding the entry in supported arrays
Loading The Library
Filesets processed: 1 of 2 (Total time: 1 secs).
```

installp: APPLYING software for:
VRTSNTAPapm 5.0.0.0

. << Copyright notice for VRTSNTAPapm >>
Copyright ? 1990-2006 Symantec Corporation. All rights reserved.

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

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

. << End of copyright notice for VRTSNTAPapm >>.

Finished processing all filesets. (Total time: 3 secs).

```
+-----+
|                               |
|                               | Summaries: |
|                               |
+-----+
```

Installation Summary

Name	Level	Part	Event	Result
VRTSNTAPas1	5.0.3.0	USR	APPLY	SUCCESS
VRTSNTAPas1	5.0.3.0	ROOT	APPLY	SUCCESS
VRTSNTAPapm	5.0.0.0	USR	APPLY	SUCCESS
VRTSNTAPapm	5.0.0.0	ROOT	APPLY	SUCCESS

Obtaining multipathing information

Information provided by the ASL

For NetApp storage systems, the ASL reports the array type as A/A-NETAPP to the device discovery layer (DDL). All primary (active) paths can be used for I/O concurrently. When all primary paths to a LUN are unavailable in a clustered storage configuration, the secondary (standby) paths are used.

For additional information, see the following documents:

- ◆ *Veritas Array Support Library Developer's Guide* for details about load balancing.
- ◆ *Veritas Volume Manager Administrator's Guide* for details about system management.

How the storage system's FC failover mode corresponds to array types

The ASL also supports direct attached, non-clustered configurations, including NearStore models. These configurations have no cfmodes settings. ASL reports these configurations as Active/Active (A/A) array types. For more details about array types, see the *Veritas Array Support Library Developer's Guide*.

How sanlun displays the array type

The Host Utilities SAN Toolkit includes the `sanlun` utility in the `/opt/NetApp/santools/bin` directory. You use the `sanlun` utility to display information about paths to LUNs on the storage system. When the ASL is installed and the LUN is controlled by VxVM, the output of the `sanlun` command displays the `Multipath_Policy` as A/A-NETAPP.

Using VxVM to display available paths

To obtain information about available paths to a LUN using VxVM, complete the following steps.

Step	Action
1	Enter the following command to view all devices: <pre>vxdisk list</pre> Result: The VxVM management interface displays the vxdisk device, type, disk, group and status. It also shows which disks are managed by VxVM. Example: The following is an example of the type of output you see with the vxdisk list command. <pre># vxdisk list DEVICE TYPE DISK GROUP STATUS FAS9600_0 auto:cdsdisk FAS9600_1 VRTS_CFS online shared FAS9600_1 auto:cdsdisk FAS9600_7 VRTS_CFS online shared FAS9600_2 auto:cdsdisk FAS9600_4 VRTS_CFS online shared FAS9600_3 auto:cdsdisk FAS9600_3 VRTS_CFS online shared FAS9600_4 auto:cdsdisk FAS9600_5 VRTS_CFS online shared FAS9600_5 auto:cdsdisk FAS9600_6 VRTS_CFS online shared FAS9600_6 auto:cdsdisk FAS9600_8 VRTS_CFS online shared FAS9600_7 auto:cdsdisk FAS9600_10 VRTS_CFS online shared FAS9600_8 auto:cdsdisk FAS9600_11 VRTS_CFS online shared FAS9600_9 auto:cdsdisk FAS9600_2 VRTS_CFS online shared (The output has been truncated to make the document easier to read.)</pre>

Step	Action																																																																																																																														
2	On the host console, enter the following command to display path information for the device you want: <pre>vxdmpadm getsubpaths dmpnodename=<device></pre> where <device> is the name listed under the output of the <code>vxdisk list</code> command. Example: The following is the type of output you should see. <pre># vxdmpadm getsubpaths dmpnodename=FAS30200_29</pre> <table><tr><th>NAME</th><th>STATE [A]</th><th>PATH-TYPE [M]</th><th>CTLR-NAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>hdisk3</td><td>ENABLED (A)</td><td>-</td><td>fcscsi0</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk33</td><td>ENABLED (A)</td><td>-</td><td>fcscsi0</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk183</td><td>ENABLED (A)</td><td>-</td><td>fcscsi1</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk213</td><td>ENABLED (A)</td><td>-</td><td>fcscsi1</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk363</td><td>ENABLED (A)</td><td>-</td><td>fcscsi2</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk393</td><td>ENABLED (A)</td><td>-</td><td>fcscsi2</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk543</td><td>ENABLED (A)</td><td>-</td><td>fcscsi3</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk573</td><td>ENABLED (A)</td><td>-</td><td>fcscsi3</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk753</td><td>ENABLED (A)</td><td>-</td><td>fcscsi0</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk783</td><td>ENABLED (A)</td><td>-</td><td>fcscsi0</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk1085</td><td>ENABLED (A)</td><td>-</td><td>fcscsi1</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk1115</td><td>ENABLED (A)</td><td>-</td><td>fcscsi1</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk1415</td><td>ENABLED (A)</td><td>-</td><td>fcscsi2</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk1445</td><td>ENABLED (A)</td><td>-</td><td>fcscsi2</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk1745</td><td>ENABLED (A)</td><td>-</td><td>fcscsi3</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk1775</td><td>ENABLED (A)</td><td>-</td><td>fcscsi3</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr></table> <pre>#</pre>	NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====	=====	=====	=====	=====	=====	=====	hdisk3	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	-	hdisk33	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	STANDBY	hdisk183	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	-	hdisk213	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	STANDBY	hdisk363	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	-	hdisk393	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	STANDBY	hdisk543	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	-	hdisk573	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	STANDBY	hdisk753	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	-	hdisk783	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	STANDBY	hdisk1085	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	-	hdisk1115	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	STANDBY	hdisk1415	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	-	hdisk1445	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	STANDBY	hdisk1745	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	-	hdisk1775	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	STANDBY
NAME	STATE [A]	PATH-TYPE [M]	CTLR-NAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																																																																									
=====	=====	=====	=====	=====	=====	=====																																																																																																																									
hdisk3	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	-																																																																																																																									
hdisk33	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk183	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	-																																																																																																																									
hdisk213	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk363	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	-																																																																																																																									
hdisk393	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk543	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	-																																																																																																																									
hdisk573	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk753	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	-																																																																																																																									
hdisk783	ENABLED (A)	-	fcscsi0	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk1085	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	-																																																																																																																									
hdisk1115	ENABLED (A)	-	fcscsi1	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk1415	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	-																																																																																																																									
hdisk1445	ENABLED (A)	-	fcscsi2	FAS3020	FAS30200	STANDBY																																																																																																																									
hdisk1745	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	-																																																																																																																									
hdisk1775	ENABLED (A)	-	fcscsi3	FAS3020	FAS30200	STANDBY																																																																																																																									

Step	Action																																																																																				
3	Enter the following command to obtain path information for a host HBA: <pre>vxdmpadm getsubpaths ctlr=<controller_name></pre> controller_name is the controller displayed under “CTRL - NAME” in the output of the vxdmpadm getsubpaths dmpnodename command you entered in Step 2. Result: The output displays information about the paths to the storage system (whether the path is a primary or secondary path). The output also lists the storage system that the device is mapped to. Example: The following is the type of output you should see. <pre># vxdmpadm getsubpaths ctlr=fscsi0</pre> <table><tr><th>NAME</th><th>STATE[A]</th><th>PATH-TYPE[M]</th><th>DMPNODENAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>hdisk9</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk55</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk101</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk147</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr></table> <pre># vxdmpadm getsubpaths ctlr=fscsi1</pre> <table><tr><th>NAME</th><th>STATE[A]</th><th>PATH-TYPE[M]</th><th>DMPNODENAME</th><th>ENCLR-TYPE</th><th>ENCLR-NAME</th><th>ATTRS</th></tr><tr><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td><td>=====</td></tr><tr><td>hdisk193</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk239</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>STANDBY</td></tr><tr><td>hdisk285</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr><tr><td>hdisk331</td><td>ENABLED(A)</td><td>-</td><td>FAS30200_86</td><td>FAS3020</td><td>FAS30200</td><td>-</td></tr></table>	NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====	=====	=====	=====	=====	=====	=====	hdisk9	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY	hdisk55	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY	hdisk101	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-	hdisk147	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-	NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS	=====	=====	=====	=====	=====	=====	=====	hdisk193	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY	hdisk239	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY	hdisk285	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-	hdisk331	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-
NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																															
=====	=====	=====	=====	=====	=====	=====																																																																															
hdisk9	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY																																																																															
hdisk55	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY																																																																															
hdisk101	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-																																																																															
hdisk147	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-																																																																															
NAME	STATE[A]	PATH-TYPE[M]	DMPNODENAME	ENCLR-TYPE	ENCLR-NAME	ATTRS																																																																															
=====	=====	=====	=====	=====	=====	=====																																																																															
hdisk193	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY																																																																															
hdisk239	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	STANDBY																																																																															
hdisk285	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-																																																																															
hdisk331	ENABLED(A)	-	FAS30200_86	FAS3020	FAS30200	-																																																																															

Using sanlun to obtain multipathing information

To obtain multipathing information by using the `san lun` command, complete the following step.

Step	Action																																																																												
1	On the host command line, enter the following command: <pre>sanlun lun show -p all</pre> Result: sanlun displays path information for each LUN. For LUNs managed by VxVM, the output displays the Veritas multipathing policy (A/A-NETAPP) and the path I/O policy. Example: The following partial output shows eight paths to one LUN managed by VxVM. <pre># sanlun lun show -p all fas3070-aix06:/vol/VRTS_SFCFS_9/VRTS_SFCFS_LUN_19 (LUN 19) DMP_NODE: fas30700_19 5g (5368709120) lun state: GOOD Filer_CF_State: Cluster Enabled Multipath_Policy: Unknown Multipath-provider: VERITAS DMP</pre> <table><tr><th>host</th><th>filer</th><th></th><th>primary</th><th>partner</th><th>Veritas</th></tr><tr><th>path</th><th>path</th><th>device</th><th>host</th><th>filer</th><th>filer</th><th>path I/O</th></tr><tr><th>state</th><th>type</th><th>filename</th><th>HBA</th><th>port</th><th>port</th><th>policy</th></tr><tr><td>up</td><td>primary</td><td>hdisk321</td><td>fcs1</td><td>4a</td><td></td><td>MinimumQ</td></tr><tr><td>up</td><td>primary</td><td>hdisk301</td><td>fcs1</td><td>4b</td><td></td><td>MinimumQ</td></tr><tr><td>up</td><td>secondary</td><td>hdisk281</td><td>fcs1</td><td></td><td>4b</td><td>MinimumQ</td></tr><tr><td>up</td><td>secondary</td><td>hdisk261</td><td>fcs1</td><td></td><td>4a</td><td>MinimumQ</td></tr><tr><td>up</td><td>primary</td><td>hdisk161</td><td>fcs0</td><td>4a</td><td></td><td>MinimumQ</td></tr><tr><td>up</td><td>primary</td><td>hdisk141</td><td>fcs0</td><td>4b</td><td></td><td>MinimumQ</td></tr><tr><td>up</td><td>secondary</td><td>hdisk121</td><td>fcs0</td><td></td><td>4b</td><td>MinimumQ</td></tr><tr><td>up</td><td>secondary</td><td>hdisk101</td><td>fcs0</td><td></td><td>4a</td><td>MinimumQ</td></tr></table> <pre>#</pre>	host	filer		primary	partner	Veritas	path	path	device	host	filer	filer	path I/O	state	type	filename	HBA	port	port	policy	up	primary	hdisk321	fcs1	4a		MinimumQ	up	primary	hdisk301	fcs1	4b		MinimumQ	up	secondary	hdisk281	fcs1		4b	MinimumQ	up	secondary	hdisk261	fcs1		4a	MinimumQ	up	primary	hdisk161	fcs0	4a		MinimumQ	up	primary	hdisk141	fcs0	4b		MinimumQ	up	secondary	hdisk121	fcs0		4b	MinimumQ	up	secondary	hdisk101	fcs0		4a	MinimumQ
host	filer		primary	partner	Veritas																																																																								
path	path	device	host	filer	filer	path I/O																																																																							
state	type	filename	HBA	port	port	policy																																																																							
up	primary	hdisk321	fcs1	4a		MinimumQ																																																																							
up	primary	hdisk301	fcs1	4b		MinimumQ																																																																							
up	secondary	hdisk281	fcs1		4b	MinimumQ																																																																							
up	secondary	hdisk261	fcs1		4a	MinimumQ																																																																							
up	primary	hdisk161	fcs0	4a		MinimumQ																																																																							
up	primary	hdisk141	fcs0	4b		MinimumQ																																																																							
up	secondary	hdisk121	fcs0		4b	MinimumQ																																																																							
up	secondary	hdisk101	fcs0		4a	MinimumQ																																																																							

Optimizing the Veritas DMP Restore Daemon

The Veritas DMP restore daemon requirements

You should set the restore policy of the Veritas DMP restore daemon and specify a polling interval based on the recommended values. These settings determine the frequency that the Veritas daemon checks paths between the host and the storage system. By default, the restore daemon checks disabled paths every 300 seconds.

For the most up-to-date recommendations on the values, see the *General Operational Considerations* page of the *FC AIX Host Utilities Release Notes*. The *Release Notes* are online in the SAN/iSAN Information Library at <http://now.netapp.com/NOW/knowledge/docs/san/>.

Setting the restore daemon interval

You must set the restore daemon interval value to 60 seconds. Doing this improves I/O failover handling.

To set the vxrestore daemon, complete the following steps.

Step	Action
1	Stop the restore daemon. <code>/usr/sbin/vxdmpadm stop restore</code>
2	Change the restore daemon setting to 60. <code>/usr/sbin/vxdmpadm start restore policy=check_disabled interval=60</code> Note — This step reconfigures and restarts the restore daemon without the need for an immediate reboot.

Step	Action
3	Veritas Storage Foundation 5.0 MP1 and earlier: If you are using one of these versions of Veritas Storage Foundation, perform these steps to ensure that new restore daemon interval persists across reboots. - Edit the <code>/etc/init.d/vxvm-sysboot</code> file to set the interval to 60 seconds. <code>restore_daemon_opts="interval=60 policy=check_disabled"</code> By default, the restore daemon options are: <code>restore_daemon_opts="interval=300 policy=check_disabled"</code> - Save and exit the <code>/etc/init.d/vxvm-sysboot</code> file. For Veritas Storage Foundation 5.0 MP3: Enter the following command to change the restore daemon interval to 60 seconds: <code>vxddmpadm settune dmp_restore_interval=60</code> When you use the <code>vxddmpadm</code> command to make the change, it persists across reboots.
4	Verify the changes. <pre>/usr/sbin/vxddmpadm stat restored</pre> Result: The command output shows the status of the vxrestore daemon, as shown by the output example below: <pre>The number of daemons running : 1 The interval of daemon: 60 The policy of daemon: check_disabled</pre>

Understanding ASL error message types

Understanding error message severity

The ASL works silently and seamlessly with the VxVM DDL. If an error, misconfiguration, or malfunction occurs, messages from the library are logged to the console using the host’s logging facility. The following table lists the severity and definition of these messages.

If have a problem or don’t understand the error message, call Symantec Technical Support for help.

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

About this chapter

This chapter provides a high level look at the tasks involved in creating and working with LUNs. It provides pointers to chapters containing instructions on configuring and managing LUNs that are specific to one AIX environment (Veritas, Native OS, or PowerVM).

Topics in this chapter

This chapter provides information on the following topics:

- ◆ [“Overview of LUN configuration and management”](#) on page 70
- ◆ [“Creating and mapping igroups and LUNs”](#) on page 72
- ◆ [“High level methods for discovering LUNs”](#) on page 74

Overview of LUN configuration and management

Task overview

LUN configuration and management involves a number of tasks. The following table summarizes the tasks for all the supported AIX environments. If a task does not apply to all environments, then the table specifies which environments it does apply to. You only need perform the tasks that are specific to your environment.

Task	Discussion	Additional information
1. Create, configure, 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 and configure 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 iSCSI and FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml .
2. Native OS, PowerVM environments: 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 117.
3. Discover and configure LUNs	You can discover LUNs using the following commands: ◆ Veritas: <code>cfgmgr</code> followed by <code>vxdctl enable</code> ◆ Native OS <code>cfgmgr</code> ◆ PowerVM <code>cfgdev</code>	Refer to “ High level methods for discovering LUNs ” on page 74 for an overview of these tasks.

Task	Discussion	Additional information
4. (Optional) Create volume groups and file systems	You can create volume groups and file systems using the Veritas or AIX tools that you normally use for these tasks.	Veritas environments: Refer to “Working with LUNs and VxVM volumes” on page 81. Native OS environments: Refer to “(Native OS only) Creating volume groups and file systems” on page 98.
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 LUN, HBA Information ” on page 103.

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.

Note

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

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

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 detailed information about creating and managing LUNs, see the *Data ONTAP® Block Access Management Guide for iSCSI and FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

Best practices for creating igroups and LUNs

As you create the igroups and LUNs, keep the following best practices in mind:

- ◆ Disable scheduled snapshots.
- ◆ 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 the storage system's root volume. The default root volume is `/vol/vol0`.
- ◆ Ensure that the LUNs are at the root of a qtree in a volume.
- ◆ Enable space reservations on the LUNs you create.
- ◆ If you have an active/active configuration, use different LUN IDs on each storage system in the configuration.

Discovering LUNs on the host

The next task you must complete involves discovering the LUNs. LUN discovery procedures vary depending on your multipathing requirements and the software you use.

- ◆ **Veritas environments:** If your configuration uses DMP, refer to [“Discovering LUNs on a host using VxVM and DMP”](#) on page 78.

Note

Before you configure LUNs to be managed by VxVM, make sure you installed the ASL and APM according to the procedures in [“Veritas: Installing the ASL and APM”](#) on page 47.

- ◆ **Native OS environments:** If your configuration has no multipathing requirements, refer to [“Discovering LUNs with native AIX LVM \(no multipathing\)”](#) on page 86.
- ◆ **Native OS environments:** If your configuration includes multipathing, refer to [“\(Native OS only\) Discovering LUNs with MPIO software”](#) on page 88
- ◆ **Native OS, PowerVM environments:** If you have a VIO server, refer to [“\(PowerVM only\) Discovering LUNs with MPIO on VIO servers”](#) on page 92.
- ◆ **Native OS environments:** If you are unsure of your configuration and software, refer to [“\(Native OS only\) Creating volume groups and file systems”](#) on page 98.

Utility for displaying information

The `sanlun` utility, which is included in the Host Utilities, provides information about paths to your devices and how they map to LUNs on the storage system. It also displays information about the host HBAs.

For more information, see:

- ◆ [“Using sanlun to display LUN information”](#) on page 104
- ◆ [“Using sanlun to display host HBA information”](#) on page 109

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).

High level methods for discovering LUNs

Commands for discovering and configuring LUNs

The commands you use to discover and configure LUNs on the host depend on your system environment:

- ◆ Native OS environments without multipathing using Native AIX Logical Volume Manager (LVM)
- ◆ Veritas environments using VxVM and DMP
- ◆ Native OS with MPIO
- ◆ Native OS with VIO servers and MPIO

Discovering 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 “[Discovering LUNs with native AIX LVM \(no multipathing\)](#)” on page 86 for additional information.

Configuring LUNs with Veritas software (DMP multipathing)

To configure LUNs using VxVM and DMP:

1. Use `cfgmgr` to get the host to discover the new LUNs.
2. Use `vxctl enable` to allow VxVM to recognize AIX hdisks.
3. Use VxVM tools (VEA, `vxdiskadm`, CLI) or `smit` to add disks to the VxVM disk group and initialize them for use by VxVM.
4. Use VxVM tools (VEA, CLI) or `smit` to create a volume
5. Use VxVM tools (VEA, CLI) or `smit` to configure data access (raw or file system)

Refer to “[Discovering LUNs on a host using VxVM and DMP](#)” on page 78 for additional information.

Non-PowerVM hosts: configuring LUNs with native AIX LVM (MPIO multipathing)

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

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 “(Native OS only) Discovering LUNs with MPIO software” on page 88 for additional information.

VIO servers: configuring LUNs with native AIX LVM (MPIO multipathing)

While the steps for working with LUNs on VIO servers are similar to those on hosts using MPIO without PowerVM, 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 “[\(PowerVM only\) Discovering LUNs with MPIO on VIO servers](#)” on page 92 for additional information.

About this chapter

This chapter describes how to enable a host running Veritas Storage Foundation to discover LUNs on the storage system. It also describes how to manage LUNs on hosts utilizing Veritas and the Veritas Volume Manager (VxVM).

Note

For general information on working with LUNs, see “[Basic Tasks for Working with LUNs](#)” on page 69.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Discovering LUNs on a host using VxVM and DMP](#)” on page 78
- ◆ “[Working with LUNs and VxVM volumes](#)” on page 81

Discovering LUNs on a host using VxVM and DMP

Discovering the new LUNs

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

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

Step	Action
1	Log in as root on the host.
2	Execute the <code>cfgmgr</code> command on the AIX host to discover the newly created LUNs: <pre># cfgmgr</pre> Result: The <code>cfgmgr</code> command causes initiators to log into the fabric, the system to check for new devices, and the system to create new device entries.
3	Enter <code>lsdev -Cc disk</code> to verify that the host has discovered the LUNs. Example: This example shows the hdisks discovered by the AIX host. <pre># lsdev -Cc disk hdisk0 Available 06-08-02 NetApp LUN hdisk1 Available 06-08-02 NetApp LUN hdisk2 Available 06-08-02 NetApp LUN hdisk3 Available 06-08-02 NetApp LUN hdisk4 Available 06-08-02 NetApp LUN hdisk5 Available 06-08-02 NetApp LUN hdisk6 Available 06-08-02 NetApp LUN hdisk7 Available 06-08-02 NetApp LUN hdisk8 Available 06-08-02 NetApp LUN hdisk9 Available 06-08-02 NetApp LUN hdisk10 Available 06-08-02 NetApp LUN hdisk11 Available 06-08-02 NetApp LUN hdisk12 Available 06-08-02 NetApp LUN #</pre>

Step	Action
4	Enter <code>vxdtl enable</code> to enable VxVM to recognize each known hdisk. Result: The <code>vxdtl</code> utility recognizes and assigns names to LUNs on the storage system. Example: The following uses <code>vxdisk list</code> command to show the results of <code>vxdtl enable</code>. <pre>#-> vxdisk list fas30700_19 Device: fas30700_19 devicetag: fas30700_19 type: auto hostid: disk: name= id=1225223596.367.p505-aix01 group: name=CoorDG id=1225223598.369.p505-aix01 info: format=cddisk,privoffset=256 flags: online ready private autoconfig pubpaths: block=/dev/vx/dmp/fas30700_19 char=/dev/vx/rdmp/fas30700_19 guid: {0ef25744-1dd2-11b2-9165-1ee0d5542ac3} udid: NETAPP%5FLUN%5F30012325%5FHn%2FOh4KwroEq site: - version: 3.1 iosize: min=512 (bytes) max=512 (blocks) public: slice=0 offset=65792 len=10415872 disk_offset=0 private: slice=0 offset=256 len=65536 disk_offset=0 update: time=1225223607 seqno=0.9 ssb: actual_seqno=0.0 headers: 0 240 configs: count=1 len=48144 logs: count=1 len=7296 Defined regions: config priv 000048-000239[000192]: copy=01 offset=000000 enabled config priv 000256-048207[047952]: copy=01 offset=000192 enabled log priv 048208-055503[007296]: copy=01 offset=000000 enabled lockrgn priv 055504-055647[000144]: part=00 offset=000000 Multipathing information: numpaths: 8 hdisk261 state=enabled hdisk281 state=enabled hdisk101 state=enabled hdisk121 state=enabled hdisk161 state=enabled hdisk141 state=enabled hdisk321 state=enabled hdisk301 state=enabled</pre>

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

Working with LUNs and VxVM volumes

Adding the disk to a disk group

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

Adding a disk to a disk group involves specifying how you want to configure the disks. Be sure to have answers to the following questions before you begin:

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

There are several methods for adding a disk to a disk group. Complete the steps for the method you prefer. Refer to Veritas documentation for additional information.

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

Creating a volume

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

Complete the steps for the method you prefer. Refer to Veritas documentation for additional information.

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

Access storage on the Veritas volume

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

Complete the steps for the method you prefer. Refer to Veritas documentation for additional information.

Method	Action
VEA	1. Select the volume that will contain the file system. 2. Specify Actions, File System, and New File System. 3. Specify the file system requirements.

Method	Action
CLI	◆ To create the file system, enter: <code>mkfs -V vxfs /dev/vx/rdisk/diskgroup/volume</code> ◆ To mount the file system, enter: <code>mount -V vxfs /dev/vx/dsk/diskgroup/volume mount_point</code>
smit	Enter the <code>smit crvxfslvstd</code> command.

Increasing the size of a LUN

If you are using Veritas Storage Foundation 5.0, you can use Dynamic LUN Expansion (DLE) to increase the size of the LUN. You can use the `vxdisk resize` command to automatically update the disk information.

On the storage system, you must run the `lun resize` command.

The only requirement for the DLE operation is that the disk group have at least two disks. This is because the DLE operation temporarily removes the disk from the disk group. It is not possible to remove the last disk from the disk group, so the operation fails if there is only one disk group.

For more information on DLE, see the Veritas documentation. For information on increasing the size of LUNs, see the *Data ONTAP® Block Access Management Guide for iSCSI and FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.

About this chapter

This chapter describes how to make LUNs visible to the host that is using a Native OS or PowerVM environment. It also describes performing tasks such as creating volume groups and file systems.

Note

For general information on working with LUNs, see “[Basic Tasks for Working with LUNs](#)” on page 69.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[Discovering LUNs with native AIX LVM \(no multipathing\)](#)” on page 86
- ◆ “[\(Native OS only\) Discovering LUNs with MPIO software](#)” on page 88
- ◆ “[\(PowerVM only\) Discovering LUNs with MPIO on VIO servers](#)” on page 92
- ◆ “[\(Native OS only\) Creating volume groups and file systems](#)” on page 98

Discovering LUNs with native AIX LVM (no multipathing)

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: # cfgmgr Result: The <code>cfgmgr</code> command causes the AIX system to complete the follow tasks: ◆ 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>

(Native OS only) Discovering LUNs with MPIO software

Getting a host using MPIO to discover the new LUNs

If you are using the Native OS environment with MPIO, you can get a host to discover the new LUNs, complete the following steps.

These steps do not apply to a host that is running PowerVM. For those steps, see “(PowerVM only) Discovering LUNs with MPIO on VIO servers” on page 92.

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 complete the following tasks: ◆ 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 to use when you perform the path configuration in Step 5. 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 0x3f000000000000 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, you must have ALUA set on the storage system. ALUA automatically sets the paths. To determine if ALUA is enabled, see “Check the ALUA settings on the storage system” on page 117. ◆ 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 following format: <code>dotpaths [-hqv] [hdiskN ...]</code> You can use the command line options for the following tasks: ❖ 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 sets the priority only 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.

(PowerVM only) Discovering LUNs with MPIO 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.
2	Execute the <code>cfgdev</code> command on the AIX host to discover the newly created LUNs. Result: This command causes the VIO server to complete the following tasks: ◆ 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 <code>cfgdev</code> command to discover the LUNs. <pre>vioserver\$ cfgdev</pre>

Step	Action
3	Enter the <code>lsdev -type disk</code> command 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> 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	Execute the <code>lsdev</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> 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 00c0ca0e99e910ca000000000000000 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 the following command: <pre>oem_setup_env</pre> 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>
6	Order and prioritize the paths to the storage system LUNs. ◆ If you are running a version of Data ONTAP that supports ALUA, you must have ALUA set on the storage system. ALUA automatically sets the paths. To determine if ALUA is enabled, see “Check the ALUA settings on the storage system” on page 117. ◆ 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: <pre>dotpaths [-hqv] [hdiskN ...]</pre> You can use the command line options to complete the following tasks: ❖ 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 sets the priority only 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. The following is an example of using the <code>dotpaths</code> 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 hdisk9. <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$
```

(Native OS only) Creating volume groups and file systems

Creating a volume group

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

Note See the appropriate 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.																																		
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 colspan="2">[Entry Fields]</td></tr><tr><td>VOLUME GROUP name</td><td>[newvg]</td></tr><tr><td>Physical partition SIZE in megabytes</td><td>4 +</td></tr><tr><td>* PHYSICAL VOLUME names</td><td>[hdisk6] +</td></tr><tr><td>Activate volume group AUTOMATICALLY at system restart?</td><td>yes +</td></tr><tr><td>Volume Group MAJOR NUMBER</td><td>[] + #</td></tr><tr><td>Create VG Concurrent Capable?</td><td>no +</td></tr><tr><td>Auto-varyon in Concurrent Mode?</td><td>no +</td></tr><tr><td>LTG Size in kbytes</td><td>128 +</td></tr><tr><td>LTG Size in kbytes</td><td>128 =Help</td></tr><tr><td>F2=Refresh</td><td>F3=Cancel</td></tr><tr><td>F5=Reset</td><td>F6=Command</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td></tr><tr><td>F4=List</td><td>F7=Edit</td></tr><tr><td></td><td>F8=Image</td></tr><tr><td>Enter=Do</td><td></td></tr><tr><td colspan="2">-----</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	F5=Reset	F6=Command	F9=Shell	F10=Exit	F4=List	F7=Edit		F8=Image	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																																		
F5=Reset	F6=Command																																		
F9=Shell	F10=Exit																																		
F4=List	F7=Edit																																		
	F8=Image																																		
Enter=Do																																			

Step	Action
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 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.
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>

Accessing 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.

Creating a file system

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

Step	Action
1	To use SMIT to set up a file system, enter the following command: smit fs
2	Enter the information to set up the file system at the SMIT Add a Journaled File System screen. ----- -->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 +-----+

Step	Action																																													
3	Select the volume group you created in the section “Creating a volume group” on page 98. ----- Add a Standard Journaled File System Type or select values in entry fields. Press Enter AFTER making all desired changes. <table><tr><td></td><td>[Entry Fields]</td><td></td></tr><tr><td>Volume group name</td><td>newvg</td><td></td></tr><tr><td>* SIZE of file system (in 512-byte blocks)</td><td>[2140000]</td><td>#</td></tr><tr><td>* MOUNT POINT</td><td>[/newfs]</td><td></td></tr><tr><td>Mount AUTOMATICALLY at system restart?</td><td>yes</td><td>+</td></tr><tr><td>PERMISSIONS</td><td>read/write</td><td>+</td></tr><tr><td>Mount OPTIONS</td><td>[]</td><td>+</td></tr><tr><td>Start Disk Accounting?</td><td>no</td><td>+</td></tr><tr><td>Fragment Size (bytes)</td><td>4096</td><td>+</td></tr><tr><td>Number of bytes per inode</td><td>4096</td><td>+</td></tr><tr><td>Allocation Group Size (MBytes)</td><td>8</td><td>+</td></tr></table> <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F4=List</td></tr><tr><td>F5=Reset</td><td>F6=Command</td><td>F7=Edit</td><td>F8=Image</td></tr><tr><td>F9=Shell</td><td>F10=Exit</td><td>Enter=Do</td><td></td></tr></table> -----		[Entry Fields]		Volume group name	newvg		* SIZE of file system (in 512-byte blocks)	[2140000]	#	* MOUNT POINT	[/newfs]		Mount AUTOMATICALLY at system restart?	yes	+	PERMISSIONS	read/write	+	Mount OPTIONS	[]	+	Start Disk Accounting?	no	+	Fragment Size (bytes)	4096	+	Number of bytes per inode	4096	+	Allocation Group Size (MBytes)	8	+	F1=Help	F2=Refresh	F3=Cancel	F4=List	F5=Reset	F6=Command	F7=Edit	F8=Image	F9=Shell	F10=Exit	Enter=Do	
	[Entry Fields]																																													
Volume group name	newvg																																													
* SIZE of file system (in 512-byte blocks)	[2140000]	#																																												
* MOUNT POINT	[/newfs]																																													
Mount AUTOMATICALLY at system restart?	yes	+																																												
PERMISSIONS	read/write	+																																												
Mount OPTIONS	[]	+																																												
Start Disk Accounting?	no	+																																												
Fragment Size (bytes)	4096	+																																												
Number of bytes per inode	4096	+																																												
Allocation Group Size (MBytes)	8	+																																												
F1=Help	F2=Refresh	F3=Cancel	F4=List																																											
F5=Reset	F6=Command	F7=Edit	F8=Image																																											
F9=Shell	F10=Exit	Enter=Do																																												
4	Press Enter. ----- COMMAND STATUS Command: OK stdout: yes stderr: no Before command completion, additional instructions may appear below. Based on the parameters chosen, the new /newfs JFS file system is limited to a maximum size of 134217728 (512 byte blocks) New File System size is 2162688 <table><tr><td>F1=Help</td><td>F2=Refresh</td><td>F3=Cancel</td><td>F6=Command</td></tr><tr><td>F8=Image</td><td>F9=Shell</td><td>F10=Exit</td><td>/=Find</td></tr><tr><td>n=Find Next</td><td></td><td></td><td></td></tr></table> -----	F1=Help	F2=Refresh	F3=Cancel	F6=Command	F8=Image	F9=Shell	F10=Exit	/=Find	n=Find Next																																				
F1=Help	F2=Refresh	F3=Cancel	F6=Command																																											
F8=Image	F9=Shell	F10=Exit	/=Find																																											
n=Find Next																																														

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

About this chapter

This chapter provides information on how to use the `sanlun` utility to display information about LUNs and host HBAs.

The `sanlun` utility provides information about paths to your devices and how they map to LUNs on the storage system. You can also use this tool to display information about the host HBAs.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Using sanlun to display LUN information”](#) on page 104
- ◆ [“Using sanlun to display host HBA information”](#) on page 109

Using sanlun to display LUN information

Displaying host LUN information

To use sanlun to display information about the LUNs connected to the host, complete the following steps.

Step	Action
1	Log in as root on the host and go to the /opt/NetApp/santools/bin directory.
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>storage system name</i> <i>storage system name:storage system pathname</i>] -p</pre> <code>-v</code> produces verbose output (you cannot use the <code>-v</code> on the same command line as the <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>storage system name</i> is the host name of the storage system. ◆ <i>storage system pathname</i> is the path name of the LUN on the storage system. (Multipathing environments only) <code>-p</code> provides information about the primary and secondary paths. You cannot use the <code>-d</code> option if you use <code>-p</code>. Use the following format: <pre>sanlun lun show -p [<i>storage system name:storage system pathname</i> <i>storage system name</i> <i>all</i>]</pre> 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
	For example, you might enter: ◆ sanlun lun show -p to display a listing of all the paths associated with the LUN. This information is useful if you need to set up path ordering or troubleshoot a problem with path ordering. ◆ sanlun lun show -d /dev/hdisk<x> to display the summary listing of the LUN(s) associated with the host device /dev/hdisk<x> (x is the number of the hdisk; for example, you might have hdisk1 or hdisk4). ◆ sanlun lun show -v all to display verbose output for all the LUN(s) currently available on the host. ◆ sanlun lun show toaster to display a summary listing of all the LUNs available to the host served by the storage system called toaster. ◆ sanlun lun show toaster:/vol/vol0/lun0 to display a summary listing of all the LUNs available to the host served by lun0 on toaster. Note_____ When you specify either the <code>sanlun lun show <storage_system_name></code> or the <code>sanlun lun show <storage_system_name:storage_system_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. _____

Interpreting sanlun lun show -p output

The `sanlun lun show -p` command provides the following information:

- ◆ path state—Whether the path is enabled or disabled.
- ◆ path type—Whether the path to the LUN is primary, secondary, or standby.
 - ❖ Primary paths communicate directly using the adapter on the local storage system.
 - ❖ Secondary paths are proxied to the partner storage system over the configuration interconnect.
- ◆ device filename—The special device file name for the disk on AIX that represents the LUN.
- ◆ host HBA—The name of the initiator HBA on the host.
- ◆ local storage system port—The port that provides direct access to a LUN. This is always a primary path.

- ◆ partner storage system port—The port that provides passive path failover. This is always a secondary path.

After the failover of a storage system, the `sanlun lun show -p` command reports secondary paths as secondary but enabled, because these are now the active paths.

- ◆ multipathing policy—The multipathing policy is one of the following:
 - ❖ A/A (Active/Active)—DMP uses more than one path concurrently for traffic. The A/A policy indicates that the B ports on the storage system HBAs are in standby mode and become active only during a takeover. This means that there are two active paths to the configuration at any given time.
 - ❖ A/A-NETAPP—Groups of LUNs connected to a single controller failover as a single failover entity. Failover occurs at the controller level and not at the LUN level.
 - ❖ The primary and secondary controllers are each connected to a separate group of LUNS. If a single LUN in the primary controller's LUN group fails, all LUNs in that group fail over to the secondary controller's passive LUN group. This array type allows load balancing of I/O across multiple primary paths. Refer to Veritas documentation for details about load balancing.

Examples of using `sanlun lun show -p`

The following are examples of using the `sanlun lun show -p` command and the type of output it displays.

Veritas example: This example displays `sanlun lun show -p` output for a system with Veritas multipathing.

```
# sanlun lun show -p
tester:/vol/norm_1_2/norm_vx-aix01_lun6 (LUN 8)      DMP_NODE: FAS30200_80
  10g (10737418240)    lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy:  A/A-NETAPP
Multipath-provider:  VERITAS DMP
```

host path state	filer path type	device filename	host HBA	primary filer port	partner filer port	Veritas path I/O policy
up	primary	hdisk95	fcs0	0d		MinimumQ
up	primary	hdisk9	fcs0	0c		MinimumQ
up	secondary	hdisk611	fcs1		0b	MinimumQ
up	secondary	hdisk525	fcs1		0a	MinimumQ
up	secondary	hdisk439	fcs0		0d	MinimumQ
up	secondary	hdisk353	fcs0		0c	MinimumQ
up	primary	hdisk267	fcs1	0b		MinimumQ
up	primary	hdisk181	fcs1	0a		MinimumQ

```
tester:/vol/nate_1_2/norm_vx-aix01_lun4 (LUN 7)      DMP_NODE: FAS30200_79
  10g (10737418240)    lun state: GOOD
Filer_CF_State: Cluster Enabled  Multipath_Policy:  A/A-NETAPP
Multipath-provider:  VERITAS DMP
```

host path state	filer path type	device filename	host HBA	primary filer port	partner filer port	Veritas path I/O policy
up	primary	hdisk94	fcs0	0d		MinimumQ
up	primary	hdisk8	fcs0	0c		MinimumQ
up	secondary	hdisk610	fcs1		0b	MinimumQ
up	secondary	hdisk524	fcs1		0a	MinimumQ
up	secondary	hdisk438	fcs0		0d	MinimumQ
up	secondary	hdisk352	fcs0		0c	MinimumQ
up	primary	hdisk266	fcs1	0b		MinimumQ
up	primary	hdisk180	fcs1	0a		MinimumQ

MPIO example: This abbreviated example displays the type of output you might see if you executed the `sanlun lun show` command with the `-p` 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 Controller target HBA port column always refers to the port on the storage system on which the LUN resides. It always represents a primary path.

```
# 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
```

MPIO path status	Controller path type	AIX MPIO path	host HBA	Controller target HBA 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

Using sanlun to display host HBA information

Displaying host HBA information

To use `sanlun` to display 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: <pre>sanlun fcp show adapter [-c [-v] [adapter name all]]</pre> <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 FC adapters. <i>adapter name</i> lists adapter information for the adapter you specify. Example 1: The following example displays the type of output you see with the <code>sanlun fcp show adapter</code> command line. This example simply displays information on the adapters. The next one includes the <code>-v</code> option, which provides more information on the adapter. <pre># sanlun fcp show adapter fcs0 WWPN:10000000abcde123 fcs1 WWPN:10000000abcde456</pre>

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

About this chapter

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

Topics covered in this chapter

This chapter discusses the following topics:

- ◆ [“Troubleshooting checklists”](#) on page 112
- ◆ [“Confirming that your configuration is set up correctly”](#) on page 114
- ◆ [“Installation fails because system does not meet requirements”](#) on page 120
- ◆ [“Tools you can use to troubleshoot problems”](#) on page 121
- ◆ [“SAN Toolkit diagnostic utilities”](#) on page 123
- ◆ [“Changing HBA attributes such as queue depth”](#) on page 130
- ◆ [“Changing the HBA queue depth with no LUNs mapped”](#) on page 135
- ◆ [“Changing the value for the queue depths on the disk”](#) on page 136
- ◆ [“General troubleshooting”](#) on page 137

Troubleshooting checklists

Check the Release Notes for the most up-to-date information

Check the *Release Notes* for a list of known problems and limitations. The *Release Notes* also contain information on how to look up information about known bugs.

You can access the *Release Notes* from the NOW site. A link to the *Release Notes* is online in the SAN/ISAN library at <http://now.netapp.com/NOW/knowledge/docs/san/>. A link is also available on the NOW Description Page for the AIX Host 5.0 Utilities.

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:

- ◆ Switch setup
- ◆ Storage system setup
- ◆ Host setup
- ◆ Storage system connectivity to the switch
- ◆ FC connectivity to the host

Storage system checklist

The following is a checklist of items that pertain to storage system setup. Make sure that the following situations are true:

- ◆ The storage system FC cfmodes is correctly set for your configuration.
- ◆ FC is running.

You can check the status of your FC service with the `fcps status` command.

- ◆ All the target ports are in a good state.
- ◆ The igroups have the ostype set to AIX.
- ◆ The LUNs are:
 - ❖ mapped to host igroups.

- ❖ in a good state as indicated by the output from the `sanlun lun show -p` command or the `lspath` command.
- ❖ using 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 FC 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`.
- ◆ **Native OS, VIO environments:** 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.
- ◆ **Native OS, VIO environments:** Use MPIO commands to check your path state and device availability.
- ◆ **Veritas environments:** Use Veritas commands to check your path state and device availability.

Confirming that your configuration is set up correctly

About the sections that follow

If you have any problems with the Host Utilities, you should make sure your configuration 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 commands, see the man pages and appropriate operating system documentation.
For more information on the Data ONTAP commands, see the Data ONTAP documentation, in particular, the *Data ONTAP® Block Access Management Guide for iSCSI and FCP* at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
- ◆ 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 `sanlun` 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 FC and iSCSI Products* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. Support for new components is added on an ongoing basis. This online information contains a complete list of supported HBAs, platforms, applications, and drivers. It is updated on a regular basis to ensure the timeliness of component information.

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 are 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
5.3.0.0
```

Check the version of the prerequisite filesets

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

```
lsllpp -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`.

```
# lsllpp -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 sisscsia0 Defined 04-08 PCI-X Dual Channel Ultra320 SCSI Adapter sisscsia1 Available 08-08 PCI-X Dual Channel Ultra320 SCSI Adapter usbhc0 Available 02-08 USB Host Controller (33103500) usbhc1 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 FC protocol on the storage system.
5. Created working volumes and qtrees (if desired) on the storage system.

Note

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

Refer to “*Where to go for more information*” on page 17 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 iSCSI and 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.

Your configuration should have no more than eight paths per LUN. For more information about zoning, see the *Configuration Guide for NetApp FC 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:

1. Configured Fibre Channel switches

It can take several minutes for the switches to boot.

2. Disk shelves

3. Storage systems

4. 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, issue a command from the host to the storage system.

One way to do this is to use a remote login, such as telnet, to connect with the storage system. Once you have logged in, enter a command at the storage system prompt.

Installation fails because system does not meet requirements

Installation checks for system requirements

The installation process 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 _____ Uninstalling software is not necessary. _____

Tools you can use to troubleshoot problems

Commands and utilities

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

- ◆ The SAN Toolkit `sanlun` utility
- ◆ The AIX `lsdev` command
- ◆ **Veritas environments:** The Veritas commands `vxddmpadm` and `vxdisk list`

Using `sanlun`

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

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

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

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

Using the `lsdev` command

The `lsdev -Cc disk` command provides information about the devices and LUNs in your configuration.

Veritas example: The following example shows the type of output displayed when you run this command in a Veritas environment. (This example has been shortened.)

```
# lsdev -Cc disk
hdisk0      Available 06-08-02      NetApp LUN
hdisk1      Available 06-08-02      NetApp LUN
hdisk2      Available 06-08-02      NetApp LUN
hdisk3      Available 06-08-02      NetApp LUN
hdisk4      Available 06-08-02      NetApp LUN
hdisk5      Available 06-08-02      NetApp LUN
hdisk6      Available 06-08-02      NetApp LUN
```

```

hdisk7      Available 06-08-02      NetApp LUN
hdisk8      Available 06-08-02      NetApp LUN
hdisk9      Available 06-08-02      NetApp LUN
hdisk10     Available 06-08-02      NetApp LUN
hdisk11     Available 06-08-02      NetApp LUN
hdisk12     Available 06-08-02      NetApp LUN
#

```

Native OS example: The following example shows the type of output displayed when you run this command in a Native OS environment. (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 that NetApp Customer support 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 that you can send to technical support.

You only use these utilities when requested to do so by NetApp’s Customer Support. The utilities include the following:

- ◆ `controller_info` and `filer_info`. See “[Collecting storage system information](#)” on page 124 for more information.

Note _____
The `controller_info` utility uses ZAPI while the `filer_info` utility uses rsh.

- ◆ `aix_info`. See “[Using the AIX data-collection tool](#)” on page 125.
- ◆ `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 125.
- ◆ `switch_info`, where *switch* is the brand of switch you are using (for example, `brocade_info`). See “[Collecting switch information](#)” on page 127.

General steps for executing the programs Although some of the scripts have different arguments that you might 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>

Step	Action
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 `controller_info` or `filer_info` to collect storage system configuration information. 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>: <pre>cd /opt/NetApp/santools/bin</pre>

Step	Action
2	At the host command prompt, enter the following command: <pre>controller_info [-d path_and_file_name] [-l user] -n <storage_system_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/netapp_filer_<storage_system_name></code>. The <code>-l user</code> option allows the choice of a user name and password on the switch. The target of the command is specified by <code>-n <storage_system_name></code>, where <i>storage_system_name</i> is the name of IP address of the storage system. Result: The output of various status commands are placed in the e output directory named <code>netapp_filer_<storage_system_name></code>, where <i>storage_system_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_<storage_system_name>.tar.gz</code>. Example: The following command line checks system_1, logging on to it as the user root.: <pre>host1> controller_info -d /opt/NetApp/santools -l root -n system_1</pre> The information from system_1 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

Veritas and Native OS environments: If you are having a problem, Customer Support might ask you to run `aix_info` to collect configuration information on your AIX system.

Note

PowerVM environments: 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 <code>aix_info</code> creates. The default location is <code>/tmp/netapp</code>. -n <file_name> specifies the name for the output file and directory names. The default pathname is <code>/tmp/netapp/ntap_aix_info.tar.gz</code>. Note Ensure that you have at least 64MB of free storage in the output directory Result: When you run <code>aix_info</code>, the utility gathers information about your AIX configuration and places it in a tarred, gzipped file. It also creates a directory and populates it with files containing output from numerous commands run on the AIX system and copies of logs such as the output of the <code>errpt</code> command. Example: Depending on the options you specify, you see output similar to the following when you run this utility: <pre># aix_info 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 of the utility name would be `mcdata_info`). This utility collects switch configuration information. You should only run this utility when asked 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" and 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.

Changing HBA attributes such as queue depth

Methods for changing the host HBA attributes

If you want to change information, such as the queue depth on the host HBA, you must change the host 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 manually 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 “[Changing the HBA queue depth with no LUNs mapped](#)” on page 135).

Changing the host HBA queue depth and rebooting

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

Step	Action
1	Enter the following command: <pre># chdev -P -l adapter_name -a num_cmd_elems=val</pre> <i>adapter_name</i> is the name of the HBA whose attributes 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 FC HBA device named fcs1 to 400. <pre># chdev -P -l fcs1 -a num_cmd_elems=400</pre>
2	Reboot the host.

Changing the host HBA queue depth without rebooting

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

Note

The commands in these steps provide an example of changing the queue depth. 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 FC 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 FC device: <pre># lsdev -C -F "name parent" awk '\$2~/^fcs1\$/ { print \$1 }'</pre> fscsi1
2	Now that you have a name for the FC 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 FC 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	Unmount the file systems. ◆ Veritas environments: Unmount any file systems and deactivate all disk groups using the <code>vxvg deports</code> command. ◆ Native OS environments: Vary off all the volume groups and unmount any file systems. Veritas example: In this example, the <code>vxvg deports</code> command unmounts the file systems that contain the disk groups called <code>my_disk_groups</code> and deactivates the disk groups: <pre># vxvg deports my_disk_groups</pre> Native OS example: To vary off the volume group called <code>my_vg</code>, 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>

Step	Action
8	Unload the FC device that you discovered in step 1. Example: This example unloads the FC 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>
10	Reload the HBA device to enable the new setting. Example: This example reloads the FC device fcs1: <pre># mkdev -l fcs1 fcs1 Available</pre>
11	If necessary, reload the FC device fscsi1. Example: The following command line reloads the FC device fscsi1: <pre># mkdev -l fscsi1 fscsi1 Available</pre>
12	If necessary, reload the child devices of the FC device. Example: The following command line reloads the child device that you unloaded in step 9: <pre># mkdev -l hdisk2 hdisk2 Available # mkdev -l hdisk3 hdisk3 Available # mkdev -l hdisk6 hdisk6 Available # mkdev -l hdisk7 hdisk7 Available</pre>

Step	Action
13	Restore the volume groups (see Step 6). ◆ Veritas environments: Import the disk groups that you deported earlier. ◆ Native OS environments: Vary on any volume groups that you varied off. Veritas example: To import the disk group called my_disk_group, enter the following command: <pre># vxdg import my_disk_groups</pre> Native OS example: To vary on the volume group called my_vg, enter the following command: <pre># varyonvg my_vg</pre>
14	Remount any file systems that you unmounted.
15	Restart any applications that you disabled.

Changing the HBA queue depth with no LUNs mapped

Changing the HBA queue depth

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

Note—————
If you have LUNs mapped to the host, you must perform the steps listed in [“Changing the host HBA queue depth without rebooting”](#) on page 131.

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

Changing 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 connectors and reconnect them. Check for faulty optical cabling. Also, check the Fibre Channel switch for excessive line errors.

About this chapter

This chapter describes of what you need to do to configure SAN booting on an AIX host running the Native OS. It also provides detailed procedures.

Currently, SAN booting is not supported for Veritas.

Note

NetApp qualifies solutions and components on an ongoing basis. To verify that SAN booting is supported in your configuration, see the *NetApp Interoperability Matrix Tool* at <http://now.netapp.com/matrix/mtx/login.do>.

Topics in this chapter

This chapter discusses the following topics:

- ◆ “[SAN boot overview](#)” on page 140
- ◆ “[Setting up a SAN Boot LUN](#)” on page 142
- ◆ “[Cloning the root volume group](#)” on page 149

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:

- ◆ Remove the hard drives from your servers and use the SAN for your boot needs. This eliminates 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 example uses a host running the AIX B 5L. You can apply this procedure to most configurations.

The procedure described here is just one example of creating a bootable LUN. This chapter does not describe other examples, such as using the AIX Alternate Disk Install, which allows you to install the OS while it is up and running.

Refer to the appropriate 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 of the operating system and your host and storage system are configured properly.
- ◆ You are logged in as root.
- ◆ You have checked the Release Notes to ensure that SAN boot steps have not changed since the manual was produced.

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.

Configuration overview

Configuring SAN booting involves the following tasks:

1. Configuring the igroup and SAN boot LUN.

- a. Determine the size of the boot LUN.
 - b. Create two igroups (one for the boot LUN and one for the data LUNs).
 - c. Create the boot LUN and any data LUNs.
 - d. Map the boot LUN to its igroup and the data LUNs to the other 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 be used only 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 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

- Creating 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 boot LUN

Creating a boot LUN involves several tasks that you must perform on the host and on the storage system. To create the boot LUN, complete the following steps.

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 149.

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 appropriate AIX product documentation to determine how much disk space is required.

Step	Action
2	Create two igroups. One igroup is for the boot LUN and the other is for any data LUNs you create. 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. To create an igroup, perform the following steps. (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 iSCSI and FCP</i> at http://now.netapp.com/NOW/knowledge/docs/ontap/ontap_index.shtml.
3	Create the boot LUN and map it to the igroup you created for it. If you have data LUNs, map them to the igroup you created. 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.

Step	Action
4	Use the <code>lun show</code> command to confirm that the boot LUN is mapped to the correct igroup. Example: The <code>lun show</code> command displays information about the boot LUN called <code>boothost</code>. <pre>filerA> lun show -g boothost /vol/sanboot/boothost 50g (53687091200) (r/w, online, mapped)</pre>
5	On the storage system with the boot LUN, create a portset that has only one port on a primary path to the boot LUN. Do not include any other ports. 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. portset create -f <i>PortSetName</i> <i>FilerName:PortName</i> <i>PortSetName</i> is the name of the portset you are creating. It can be up to 95 characters long. <i>FilerName:PortName</i> is the name of the FC port. Note _____ For details on creating portsets, see the <i>Data ONTAP Block Access Management Guide</i>. _____ Example: This example creates a portset called <code>sanboot</code> for the LUN. <pre>filerA> portset create -f sanboot controllerA:4a</pre>
6	Bind the portset to the igroup for the boot LUN. Use the command: igroup bind <i>iGroupName</i> <i>PortSetName</i> <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 filerA. 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 boothost (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>

Step	Action
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 appropriate AIX documentation.
8	Install the Host Utilities. For information on doing this, see “Installing the Host Utilities software packages” on page 27.
9	On the host, use the <code>cfgmgr</code> command to discover the NetApp LUNs.
10	On the host, enter the <code>bootinfo -b</code> command to identify which device is bound to the boot LUN. Example: In this example, the boot LUN is the device <code>hdisk0</code>: <pre>AIX_host> bootinfo -b hdisk0</pre>
11	Change the boot LUN's reserve policy for MPIO by entering the following command: <pre>AIX_Host> chdev -Pl hdisk0 -a algorithm=fail_over -a reserve_policy=single_path</pre> This command prevents SCSI-2 Reserve conflicts by setting the reserve policy to <code>single_path</code>. Note _____ You do not have to change the settings for data LUNs. Data LUNs can use the Host Utilities default settings. _____
12	Use the <code>lsattr</code> command to verify that the <code>reserve_policy</code> attribute is set to <code>single_path</code> and that the <code>Algorithm</code> attribute is set to <code>fail_over</code>. Example: The following command shows that both the <code>reserve_policy</code> attribute and <code>algorithm</code> attribute are correctly set: <pre>AIX_Host> # lsattr -El hdisk0 -a algorithm -a reserve_policy algorithm fail_over Algorithm True reserve_policy single_path Reserve Policy True</pre>
13	Restart the AIX server so that it boots from the boot LUN with the MPIO and SCSI-2 Reserve parameters you set in Step 11.

Step	Action
14	On the storage system, add host WWPNs to the boot LUN's igroup so that there is a maximum of four paths to the boot LUN. Example: The following example adds a WWPN to the boot LUN igroup. The <code>igroup show</code> command displays the WWPNs in the igroup and the igroup's portset: <pre> controllerA> igroup add SANboot 10:00:00:00:c9:50:5f:8c controllerA> igroup show SANboot BootLUN (FCP) (ostype: aix): 10:00:00:00:c9:40:56:23 (logged in on: 5a, 5b, vtic, 4b, 4a) 10:00:00:00:c9:50:5f:8c (logged in on: 5a, 5b, vtic, 4a, 4b) Pset: bootlun </pre>
15	Add the target port on the secondary path for the boot LUN to the igroup's port set. This port is on the partner in the active/active configuration. For example, if the primary path to the boot LUN is target port 4a on controllerA, then the secondary path is target port 4a on controllerA's partner. Example: The following example adds port 4a on both controllerA and its partner to the portset: <pre> contolleraA> portset add bootlun 4a controllerA> portset show bootlun (FCP): ports: controllerA 4a controllerB 4a igroups: SANboot </pre>
16	On the host, use the <code>cfgmgr</code> command to verify that the boot LUN has only the recommended four paths. Example: This example uses the <code>cfgmgr</code> command and then displays the paths using the <code>lspath</code> command. <pre> AIX_Host> cfgmgr AIX_Host> # lspath -l hdisk0 Enabled hdisk0 fscsi0 Enabled hdisk0 fscsi0 Enabled hdisk0 fscsi1 Enabled hdisk0 fscsi1 </pre>

Step	Action
17	Use the <code>lsdev</code> command to identify any ghost devices left over from the ODM device overlay from MPIO to MPIO NetApp. Example: The following example shows ghost devices with two different device handles: Defined and Unknown. <pre>lsdev -Cc disk hdisk0 Available 02-08-01 MPIO NetApp FCP Default PCM Disk hdisk1 Defined 02-08-01 MPIO NetApp FCP Default PCM Disk hdisk2 Defined 03-08-01 MPIO NetApp FCP Default PCM Disk hdisk3 Unknown 05-08-01 MPIO NetApp FCP Default PCM Disk hdisk4 Unknown 06-08-01 MPIO NetApp FCP Default PCM Disk hdisk5 Available 02-08-01 MPIO NetApp FCP Default PCM Disk hdisk6 Available 03-08-01 MPIO NetApp FCP Default PCM Disk hdisk7 Available 05-08-01 MPIO NetApp FCP Default PCM Disk hdisk8 Available 06-08-01 MPIO NetApp FCP Default PCM Disk</pre>
18	Remove the ghost devices using either the <code>odmdelete</code> command or the <code>rmdev</code> command, depending on whether your device handles were marked Unknown or Defined. Example (Unknown): For device handles marked as Unknown, use the <code>odmdelete</code> command: <pre>odmdelete -q name=hdisk3 -o CuDv 0518-307 odmdelete: 1 objects deleted</pre> Example (Defined): For device handles marked as Defined, use the <code>rmdev</code> command: <pre>rmdev -dl hdisk1 hdisk1 deleted</pre>
19	Use the <code>cifmgr</code> command to rescan the bus for MPIO NetApp FC devices.
20	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>AIX_Host> /usr/sbin/dotpaths -v</pre>

Step	Action
21	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 perform the following tasks: ◆ Return information on the last boot device. ◆ Invalidate the normal mode boot list ◆ Set a the new boot device to a normal boot list. Note_____ This procedure helps prevent a common boot issue where the boot list becomes filled with paths that cannot be used during the boot cycle. _____

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 <code>bootlist</code> command to create a new boot list that boots only from the primary paths of the boot LUN. Example: The following example contains commands to perform the following tasks: ◆ Create a boot list in normal mode that says to boot off logical volume hd5 on disk hdisk1 and displays information about the boot. ◆ Return information on the last boot device. ◆ Display a 1 if the IPL code in the ROS on the machine running the command is capable of booting from a specified disk. <pre>-> bootlist -m normal -o hdisk1 blv=hd5 -> bootinfo -b hdisk1 -> bootinfo -B hdisk1 1</pre>
4	Use the <code>ipl_varyon</code> 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 information

The following copyright statements and licenses apply to open source software components that are distributed with various versions of the NetApp Host Utilities software products. Your product does not necessarily use all the open source software components referred to below.

Where required, source code is published at the following location:

<ftp://ftp.netapp.com/frm-ntap/opensource/>

Copyrights and licenses

The following components are subject to the Perl Artistic License:

- ◆ FTP.pm portions copyright © 1995-2004 Graham Barr <gbarr@pobox.com>. All rights reserved.
- ◆ Telnet.pm portions copyright © 1997, 2000, 2002 Jay Rogers. All rights reserved.

Perl Artistic License:

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

“Package” refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

“Standard Version” refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder.

“Copyright Holder” is whoever is named in the copyright or copyrights for the package.

“You” is you, if you’re thinking about copying or distributing this Package.

“Reasonable copying fee” is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

“Freely Available” means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes, and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
 - Place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as ftp.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
 - Use the modified Package only within your corporation or organization.
 - Rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
 - Distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
 - Accompany the distribution with the machine-readable source of the Package with your modifications.
 - Accompany any non-standard executables with their corresponding Standard Version executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
 - Make other distribution arrangements with the Copyright Holder.
5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package.

You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggregated with this Package.
7. C or perl subroutines supplied by you and linked into this Package shall not be considered part of this Package.
8. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
9. THIS PACKAGE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The HBA API component is subject to the SNIA Public License 1.1.

Portions copyright © 2002 International Business Machines Corporation. All rights reserved.

Portions authored by Benjamin F. Kuo, Troika Networks, Inc.

Portions copyright © 2001 Microsoft Corporation. All rights reserved.

SNIA Public License 1.1:

1. Definitions.

1.1 “Commercial Use” means distribution or otherwise making the Covered Code available to a third party.

1.2 “Contributor” means each entity that creates or contributes to the creation of Modifications.

1.3 “Contributor Version” means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.4 “Covered Code” means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.5 “Electronic Distribution Mechanism” means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.6 “Executable” means Covered Code in any form other than Source Code.

1.7 “Initial Developer” means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.8 “Larger Work” means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.9 “License” means this document.

1.10 “Licensable” means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.11 “Modifications” means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

- ❖ Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.
- ❖ Any new file that contains any part of the Original Code or previous Modifications.

1.12 “Original Code” means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.13 “Patent Claims” means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.14 “Source Code” means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor’s choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.15 “You” (or “Your”) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, “You” includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, “control” means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty

percent (50%) of the outstanding shares or beneficial ownership of such entity

2. Source Code License.

2.1 The Initial Developer Grant. The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third-party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense, and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using, or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2 Contributor Grant. Subject to third-party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense, and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2)

separate from the Contributor Version; 3) for infringements caused by: i) third-party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1 Application of License. The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2 Availability of Source Code. Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3 Description of Modifications. You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4 Intellectual Property Matters.

(a) Third Party Claims. If Contributor has actual knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail

that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter.

(b) Contributor APIs. If Contributor's Modifications include an application programming interface and Contributor has actual knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations. Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5 Required Notices. You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be most likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity, or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity, or liability terms You offer.

3.6 Distribution of Executable Versions. You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligation of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a

different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability (excluding any liability arising from intellectual property claims relating to the Covered Code) incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7 Larger Works. You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation. If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.
5. Application of this License. This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.
6. Versions of the License.

6.1 New Versions. The Storage Networking Industry Association (the “SNIA”) may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2 Effect of New Versions. Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by the SNIA. No one other than the SNIA has the right to modify the terms applicable to Covered Code created under this License.

6.3 Derivative Works. If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases “Storage Networking Industry Association,” “SNIA,” or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the SNIA Public License. (Filling in the name of the Initial Developer,

Original Code, or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY. COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR, OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1 This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within a reasonable time after becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2 If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as “Participant”) alleging that:

(a) such Participant’s Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

8.3 If You assert a patent infringement claim against Participant alleging that such Participant’s Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior

to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4 In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. **LIMITATION OF LIABILITY.** UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.
10. **U.S. GOVERNMENT END USERS.** The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.
11. **MISCELLANEOUS.** This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation

which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. **RESPONSIBILITY FOR CLAIMS.** As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.
13. **MULTIPLE-LICENSED CODE.** Initial Developer may designate portions of the Covered Code as “Multiple-Licensed”. “Multiple-Licensed” means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of this License or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.
14. **ACCEPTANCE.** This License is accepted by You if You retain, use, or distribute the Covered Code for any purpose.

EXHIBIT A—The SNIA Public License.

The contents of this file are subject to the SNIA Public License Version 1.1 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

This file itself is the license.

Software distributed under the License is distributed on an “AS IS” basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Contributor(s): _____.

The Net_SSLeay.pm component is subject to the Net SSLeay License.

NetSSLeay License:

http://search.cpan.org/src/SAMPO/Net_SSLeay.pm-1.25/README License and Copying: Copyright © 1996-2002 Sampo Kellomaki. All rights reserved.

Distribution and use of this module is under the same terms as the OpenSSL package itself (i.e. free, but mandatory attribution; NO WARRANTY). Please consult LICENSE file in the root of the OpenSSL distribution. While the source distribution of this perl module does not contain Eric’s or OpenSSL’s code, if you use this module you will use OpenSSL library. Please give Eric and OpenSSL team credit (as required by their licenses). And remember, you, and

nobody else but you, are responsible for auditing this module and OpenSSL library for security problems, backdoors, and general suitability for your application.

SSLey License copyright © 1995-1998 Eric Young (eay@cryptsoft.com). All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape's SSL. This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)". The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related.
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgment: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)".

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "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 THE AUTHOR OR CONTRIBUTORS 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. The licence and distribution terms for any publicly available version or derivative of this code cannot be changed; i.e., this code cannot simply be copied and put under another distribution licence [including the GNU Public Licence.

The Netkit component is subject to the GNU GPL 2.0 license.

Copyright © 1989, 1991 Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

GNU General Public License 2.0:

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation’s software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution, and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION, AND MODIFICATION

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
 - a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
 - b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
 - c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
 - a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute, or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients'

exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement, or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this

License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found:

One line to give the program’s name and a brief idea of what it does. Copyright (c)

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w`. This is free software, and you are welcome to redistribute it under certain conditions; type `show c` for details.

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program ‘Gnomovision’ (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

The Open Hardware Certification component is subject to the Robin's Nest Software license.

Copyright © 1990-2008 Robin's Nest Software, Inc. All rights reserved.

Robin's Nest Software License:

This Software Provided By Robin's Nest Software Inc.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in the supporting documentation, and that the name of the author not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL HE BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE, OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

The Windows Installer XML (WiX) Toolset component is subject to the Common Public License.

Copyright © Microsoft Corporation. All rights reserved.

Common Public License:

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION, OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

b) in the case of each subsequent Contributor:

- i) changes to the Program, and
- ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution ‘originates’ from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor’s behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

“Contributor” means any person or entity that distributes the Program.

“Licensed Patents” mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

“Program” means the Contributions distributed in accordance with this Agreement.

“Recipient” means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import, and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the

rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third-party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

- a) it complies with the terms and conditions of this Agreement; and
- b) its license agreement:
 - i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;
 - ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental, and consequential damages, such as lost profits;
 - iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and
 - iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

- a) it must be made available under this Agreement; and
- b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor

who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor (“Commercial Contributor”) hereby agrees to defend and indemnify every other Contributor (“Indemnified Contributor”) against any losses, damages, and costs (collectively “Losses”) arising from claims, lawsuits, and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor’s responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL,

SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), 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 OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published,

Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

Index

Symbols

60, 124

A

AIX host

- example of successful installation 33
- example of successful uninstallation 45

aix_info

- configuration script 3
- executing 125

ALUA

- Asymmetric Logical Unit Access 6
- checking on igroup 117
- defined 6
- setting up after installing Host Utilities 117
- used with MPIO 6

APM

- definition 49
- fresh install 50
- installing 51
- obtaining 51
- uninstalling 55
- upgrading from previous version 51

ASL

- components 49
- definition 48
- fresh install 50
- information reported 60
- installing 51
- obtaining 51
- removing 55
- uninstalling 55
- upgrading from previous version 51
- using storage with 48

attributes

- changing and rebooting 130
- changing on HBAs 130
- changing with SMIT 130
- changing without rebooting 131

C

CD-ROM

- obtaining software 25

cfgmgr command 78, 86, 88, 92

cfmode

- changing 6

chdev command 130

checklists

- filer setup 112
- host setup 113
- installation tasks 12
- troubleshooting 112

cluster failover mode

- cfmode 6

commands

- oem_setup_env 10

configuration

- direct-attached 23

controller_info

- diagnostic script 3
- diagnostic utilities 123

conventions

- command x
- formatting x
- keyboard x

D

diagnostic utilities

- aix_info 123, 125
- controller_info 123
- executing 123
- filer_info 123, 124
- ontap_collect 123, 126
- switch_info 123, 127

direct-attached configuration

- set media type 23

disk

- adding to a disk group (Veritas) 81
- changing value of queue depths 136

disk group

- adding disks (Veritas) 81

- disk shelves
 - powering up 118
- DLPAR
 - Dynamic Logical Partitioning 6
- DMP
 - Veritas multipathing 9
- dotpaths
 - executing on VIO servers 177
 - set path priorities 3, 90, 95

E

- environments
 - comparing 8
 - MPIO 2
 - Native OS 2
 - PowerVM 2
 - Veritas 2
- equipment
 - powering up 118

F

- Fibre Channel
 - used with iSCSI 5
- file system
 - creating 100
- filer_info
 - diagnostic script 3
- files
 - names used in examples 24
 - switch_info 127
- fileset
 - check version 115

G

- GPFS 7

H

- HBAs
 - about 6
 - changing attributes 130
 - changing queue depth 130, 135
 - changing queue depth and rebooting 130

- changing queue depth without rebooting 131
- displaying host information 109
- installing 116
- verifying queue depth 39

host

- check OS version 114
- communicating with storage system 118
- creating volume group 98
- device for physical volume 91, 97
- discovering new LUNs, MPIO 88
- discovering new LUNs, native AIX LVM 86
- discovering new LUNs, VIO servers 92
- discovering new LUNs, VxVM and DMP 78
- displaying LUN information (Veritas) 104
- powering up 118
- PowerVM environment requirements 10
- setup checklist 113

Host Utilities

- completing installation for Veritas 70
- contents 2
- dotpath utility 90, 95
- overview 2
- running in PowerVM environment 10
- SAN Toolkit 3

I

igroup

- checking ALUA setting 117
- defined ix
- definition 70
- methods for creating 72

initiator

- defined ix

install command

- checks system requirements 120
- rerunning 120

installing

- methods of 27
- overview 13
- pre-installation tasks 22
- reinstalling 120
- using SMIT 27

iSCSI

- used with FC 5

L

- lsdev command 131
 - using for troubleshooting 121
- lspp command 115
- LUN ID
 - defined ix
- LUNs
 - best practices for creating 72
 - defined ix
 - discovering on VIO servers 92
 - discovering with AIX LVM 86
 - discovering with MPIO 88
 - discovering with VxVM and DMP 78
 - displaying host information (Veritas) 104
 - increasing size of (Veritas) 83
 - maximum transfer size with PowerVM 10
 - methods for creating 72
 - tasks for creating on filer 72
- LVM
 - AIX Logical Volume Manager 9

M

- maximum transfer size
 - LUNs with PowerVM 10
- media type
 - setting 23
- MPIO
 - ALUA 6
 - environment 2
 - multipathing 10
 - using on host 2

N

- Native OS
 - comparing environments 8
 - environment 2
 - environment requirements 9
 - multipathing 10
 - volume management 9
- NOW site
 - downloading software 25

O

- oem_setup_env command
 - PowerVM environment command 10
 - using 95
- ontap_collect
 - configuration script 3
 - executing 126
 - PowerVM environment
 - PowerVM
 - ontap_collect script 10
- oslevel command 114

P

- padmin
 - PowerVM 10
- path priorities
 - setting 90, 95
- physical volume
 - selecting device 91, 97
- powering up equipment 118
- PowerVM
 - comparing environments 8
 - environment 2
 - environment requirements 10
 - LUN transfer size 10
 - padmin 10
 - VIO 2
- protocol
 - Fibre Channel 5
 - iSCSI 5
 - multiple 5

Q

- queue depth
 - changing value on disk 136
 - changing value on HBAs 130, 135
 - verifying 39

S

- SAN Toolkit
 - components 3
 - san_version command 3

- sanlun utility 3
- san_version
 - command 3
- sanlun fcp show adapter command 109
- sanlun lun show command 104
- sanlun utility
 - about 73
 - displaying multipathing information 63
 - displaying path policy information 60
 - part of SAN Toolkit 3
 - sanlun fcp show adapter 109
 - syntax 105
 - using for troubleshooting 121
- scripts
 - aix_info 3
 - controller_info 3
 - diagnostic 3
 - filer_info 3
 - ontap_collect 3
- SMIT
 - unloading child devices 130
 - using to create a file system 100
 - using to create a volume group 98
 - using to install Host Utilities 27
 - using to uninstall Host Utilities 40
- software
 - downloading 25
 - installation tasks 27
 - obtaining 24
 - obtaining from CD-ROM 25
 - uncompressing 26
- storage
 - accessing on volume group 99
 - accessing with VxVM 82
- storage system
 - communicating with host 118
 - direct-attached 23
 - powering up 118
 - setup checklist 112
 - verifying setup 116
- switch_info
 - executing 127
- switches
 - collecting information 127
 - powering up 118

- verifying setup 117
- zoning 118

T

- target
 - defined ix
- terminology ix
- troubleshooting
 - commands and utilities 121
 - connectivity problem checklist 112
 - general 137
 - host setup checklist 113
 - storage setup checklist 112
 - support information 17
 - using lsdev command 121
 - using sanlun 121

U

- uninstalling
 - methods 40
 - using a command line 46
 - using SMIT 40

V

- Veritas
 - comparing environments 8
 - environment 2
 - environment requirements 9
 - multipathing 9
 - using on host 2
 - volume management 9
- VIO
 - PowerVM 2
- VIO servers
 - executing dotpaths 177
 - oem_setup_env command 10, 95
 - ontap_collect 126
 - troubleshooting 11
 - using on host 2
 - using with Host Utilities 10
- volume group
 - accessing storage 99
 - creating 98

volumes

creating with VxVM 82

VxVM

accessing storage 82

adding disks to a disk group 81

creating volumes 82

displaying available paths 61

Veritas volume management 9

Z

zoning 118

