
Part Number: E37669-01

APM 3.3

2012 December

Copyright © 2005, 2012, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

This software or hardware and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2005, 2012, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Table of Contents

Preface

Chapter 1 Introduction to APM

APM Features.	10
About APM and Hypervisors.	13
APM Architecture.	14
About the APM Control Path.	16
About the APM Data Path.	17
About Multipathing and Device Mapper Automation.	17
Supported Distributions.	19
APM Requirements.	20
Operating Limits.	21

Chapter 2 Install APM

Prepare to Install APM.	22
Supported SAN Protocols.	23
Supported FC SAN Hardware Components.	23
Supported iSCSI Software and Hardware.	30
Pre-Configure the SAN Host for Pillar Axiom Integration.	35
Verify Installation of the Device Mapper.	37
About APM Software Installation.	38
Download the APM Software.	38
Install (or Upgrade) the APM Software.	39
Configure SAN Host Access to the Pillar Axiom LUNs.	41
Partition and Format the LUN Disk (Optional).	44
About Load Balancing Configurations.	46
About Path Selection.	47
Remove the APM Software (Optional).	49

Chapter 3 APM Release Notes

New in this Release.	50
------------------------------	----

Known APM Issues.	51
Known Pillar Axiom Issues.	52
Known Operating System Issues.	55
Filesystems Marked Read-Only When Access to Boot LUN Is Lost.	55
Access to LUNs.	55
[13746418] SAN Dynamic Reconfiguration.	56
[13758951] Multipath Failed Path Errors.	58
[13764910] Non-Optimized Access Events.	58
[13761123] LUN State Not Updated.	59
[13761100] Too Many Files Open.	59
[13762235] Too Many Paths Cause Device Mapper to Hang.	59
[13762937, 13766418] LUNs Not Shown.	60
[13764931,13792540] Device Mapper Driver Not Loaded.	60
[13819519] CNAs Require Firmware Update.	61
[13959093] LUNs Not Seen.	61
[14708044] Xsigo Driver Fails to Restore Paths.	61
[14708089, 14711677] Adding or Deleting LUNs Requires Reboot.	61
Resolved APM Issues.	62
Additional Notes.	63
Mount iSCSI Filesystems.	63
Index.	64



List of Figures

Figure 1 APM interaction with a Pillar Axiom system. 15

Figure 2 Example host ports before APM installation. 41

Figure 3 Example host ports after APM installation. 41



List of Tables

Table 1 Typography to mark certain content. 8

Table 2 Oracle resources. 9

Table 3 APM for CentOS 6 features. 10

Table 4 Line and color key for APM interaction diagram. 14

Table 5 Supported hardware platforms. 19

Table 6 APM operating limits. 21

Table 7 Emulex adapter settings. 25

Table 8 QLogic adapter settings. 26

Table 9 Brocade adapter setting. 29

Table 10 Characteristics of SAN connections to host ports. 36

Table 11 Known Pillar Axiom issues. 52

Preface

Related Documentation

- *Pillar Axiom Customer Release Notes*
- *Pillar Axiom Glossary*
- *Pillar Axiom Administrator's Guide*
- *Pillar Axiom CLI Reference Guide*

Typographical Conventions

Table 1 Typography to mark certain content

Convention	Meaning
<i>italics</i>	Within normal text, words in italics indicate one of the following: <ul style="list-style-type: none">• A reference to a book title• New terms and emphasized words• Command variables
monospace	Indicates one of the following, depending on the context: <ul style="list-style-type: none">• The name of a file or the path to the file• <i>Output</i> displayed by the system on the command line
monospace (bold)	<i>Input</i> provided by an administrator on the command line.
>	Indicates a menu item or a navigation path in a GUI. For example, "Click Storage > Clone LUNs " means to click the Clone LUNs link on the Storage page in the GUI.
...	Used within an expression of a navigation path or within a cascading menu structure. The ellipsis indicates that one

Table 1 Typography to mark certain content (continued)

Convention	Meaning
	or more steps have been omitted from the path or menu structure. For example, in the Groups > Volume Groups > Actions > ... > Data Protection > Create menu structure, the ... implies that one or more menu items have been omitted.

Oracle Contacts

Table 2 Oracle resources

For help with...	Contact...
Support	http://www.oracle.com/support/index.html
Training	https://education.oracle.com
Documentation	<ul style="list-style-type: none"> • Oracle Technical Network: http://docs.oracle.com • From the Pillar Axiom Storage Services Manager (GUI): Support > Documentation • From Pillar Axiom HTTP access: http://system-name-ip/documentation.php where <i>system-name-ip</i> is the name or the public IP address of your system.
Documentation feedback	http://www.oracle.com/goto/docfeedback
Contact Oracle	http://www.oracle.com/us/corporate/contact/index.html

CHAPTER 1

Introduction to APM

APM Features

Pillar Axiom Path Manager (APM) is defined as:

Optional software installed on a SAN host to manage multiple paths to the Pillar Axiom system.

APM performs the following primary functions:

- Routes I/O to Pillar Axiom LUNs using only the best available data paths
- Shares traffic among the available paths and ensures that access to the LUNs is not interrupted if some paths fail
- Automatically configures the host into the Pillar Axiom Storage Services Manager and updates the configuration if the host information changes

Automatic configuration of host information enables the Pillar Axiom Storage Services Manager to report information about the instance of APM running on the host such as the number of working paths. In some environments, automatic configuration includes features such as load balancing.

Each APM release provides different features, and the features provided for each platform might vary. The following table describes the specific features implemented in this release.

Table 3 APM for CentOS 6 features

Feature	Benefit
Automatic data path failover	Automatically switches to the most suitable paths available after a path failure or fail back.
Automatic recognition of SAN hosts by the Pillar Axiom Storage Services Manager	Sends a description of the host to each Pilot management controller on connected Pillar Axiom systems, allowing the Pillar Axiom Storage Services Manager to create a definition for the host. This definition includes such information as the WWNs for each of the host FC ports, the

Table 3 APM for CentOS 6 features (continued)

Feature	Benefit
	IP addresses for any iSCSI ports, and the version of APM running on the host.
Call-Home log collection	<p>When a Pillar Axiom administrator uses the Pillar Axiom Storage Services Manager to collect system information (refer to the <i>Pillar Axiom Administrator's Guide</i> for details), the Pillar Axiom system sends a request to each connected APM host. The APM hosts collect useful diagnostic information and send it to the Pillar Axiom system, where it is bundled with any other requested information. The Pillar Axiom system can then transmit this information to Oracle Customer Support. The information collected from each APM host includes:</p> <ul style="list-style-type: none"> • Logs from the APM components • Configuration and status information from the operating system • System and error logs from the operating system <p>No customer data is transmitted.</p>
Support for FC connections to FC SAN Slammers	Makes connections to Pillar Axiom storage arrays over a high-speed FC network infrastructure.
Support for iSCSI connections to both FC and iSCSI SAN Slammers	<p>Makes connections to Pillar Axiom storage arrays over long distances using an IP network infrastructure.</p> <p>Note: iSCSI connections to FC SAN Slammers require iSCSI-to-FC routers.</p>
Support for Boot from SAN	Supports Boot from SAN on supported QLogic and Emulex Fibre Channel host bus adapters (HBAs).
FCoE CNAs	FC over Ethernet (FCoE) Converged Network Adapters (CNAs) are supported on the host.
Support for virtualization	APM supports the KVM hypervisor included in the Linux distribution.
Support for load balancing	Both static and round-robin load balancing options are supported.

Related concepts

- [*About APM and Hypervisors*](#)

Related references

- [*Supported Emulex FC Adapters and Drivers*](#)
- [*Supported QLogic FC Adapters and Drivers*](#)

About APM and Hypervisors

Pillar Axiom Path Manager (APM) supports the Kernel-based Virtual Machine (KVM) open source hypervisor, or virtual machine manager. Only the KVM hypervisor included with the Linux distribution is supported.

You can install APM on a CentOS 6 host configured as a KVM hypervisor host. Guest operating systems can access multipathed Pillar Axiom LUNs that are mapped to the hypervisor host in the same way that these systems access other disks. For example, these LUNs can be used to hold file-based containers, which can then be exposed to the guest operating systems, or the guest operating systems can access the LUNs as persistent storage disks.

For information on using disks with the KVM hypervisor, refer to the [Red Hat Enterprise Virtualization Host Configuration and Guest Installation Guide](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Virtualization_Host_Configuration_and_Guest_Installation_Guide/) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Virtualization_Host_Configuration_and_Guest_Installation_Guide/).

Guest operating systems can also use iSCSI software initiators to connect to Pillar Axiom systems and to access LUNs. In this case, APM maps the LUNs directly to the guest operating system. If the guest operating system has a version of APM that supports iSCSI, this version of APM can be installed on the guest and used to manage multiple iSCSI paths to the Pillar Axiom LUNs in the same way as APM would be used on a physical host.

APM Architecture

Pillar Axiom Path Manager (APM) manages the Linux multipath framework and communicates with Pillar Axiom systems on a control path, which is separate from the data path. The Linux multipath framework manages the LUN data access paths.

The following figure illustrates how the APM software installed on a SAN host interacts with a Pillar Axiom system. Refer to the table below to determine the significance of the lines and colors in the figure.

Table 4 Line and color key for APM interaction diagram






Graphic element	Description
————	Data path
-----	Control path
	Pillar Axiom hardware and software
	Non-Pillar Axiom hardware and software
	SAN host kernel space
	SAN host user space

Figure 1 APM interaction with a Pillar Axiom system



Legend	
1 User	10 Linux multipath framework
2 User application	11 iSCSI software initiator (iSCSI)
3 SAN host	12 TCP/IP driver (iSCSI)
4 APM daemon	13 HBA driver (FC) or NIC driver (iSCSI)
5 Control path (all dashed lines)	14 HBA (FC) or NIC (iSCSI)
6 Pillar Axiom administrator	15 SCSI over FC (FC) or iSCSI over IP (iSCSI)
7 Pillar Axiom CLI or GUI	16 Data path (all solid lines)
8 Encrypted XML over TCP/IP	17 Pillar Axiom system
9 Network card	18 Brick storage pool

Related concepts

- [About the APM Control Path](#)
- [About the APM Data Path](#)
- [About Multipathing and Device Mapper Automation](#)

About the APM Control Path

The Pillar Axiom Path Manager (APM) control path provides a path separate from the data path to manage multipathing and communication.

The APM software uses a daemon to control multipathing and communication. The APM daemon uses the control path to perform the following actions:

- Get FC and iSCSI port information from the HBA and converged network adapter (CNA) drivers and from the iSCSI software initiator.
- Configure the Linux multipath framework.
- Send information such as host attributes and statistics to the Pilot management controller and, on request, collect logs from the host.

The APM daemon sends a description of the host to the Pilot on each connected Pillar Axiom system. In the Pillar Axiom Storage Services Manager, this description creates a definition for the host that includes any FC ports in the host and, if iSCSI is configured, the name of the iSCSI initiator.

The graphical user interface (GUI) and command line interface (CLI) list the World Wide Names (WWNs) of the FC ports in the host and the IP addresses that are used to make iSCSI connections to the Pillar Axiom system.

If you use iSCSI on the host to connect to a FC Slammer storage controller through an iSCSI-to-FC router, these connections are described as FC connections. Within APM, the connections appear to originate from the FC ports on the switch that are assigned to the host iSCSI initiator. The WWNs of these ports are displayed as FC HBA ports on the host. The HBA model associated with these ports is reported as iSCSI-FC router.

To establish the control path to a Pillar Axiom host, that host must be able to connect to the Pillar Axiom system over the data path. As part of the connection sequence, the Slammer returns the IP address of the Pilot to the APM host over the data path. The host uses the IP address of the Pilot to establish the control path to the Pillar Axiom system.

About the APM Data Path

Pillar Axiom Path Manager (APM) uses the Linux `device-mapper` to provide paths for reading and writing data to LUNs on the Pillar Axiom system.

The Linux multipath framework is responsible for the following:

- Controls and manages all data paths to Pillar Axiom LUNs
- Groups multiple data paths to a Pillar Axiom LUN and presents this group to the host operating system as a single LUN or drive
- Determines which data paths to use
- Identifies and uses optimized data paths when possible
- Handles data path failover and failback
- Manages data path errors

See the architecture diagram for an illustration of how data flows from the host to the Pillar Axiom system.

A path that provides the best performance is referred to as an *optimized path*. It is the preferred path for data transfer.

Related references

- [APM Architecture](#)

About Multipathing and Device Mapper Automation

Pillar Axiom Path Manager (APM) uses the standard Linux 2.6 `device-mapper` functionality to provide multipathing services.

APM installs extensions to the Linux Multipath daemon that monitor the state of the Pillar Axiom LUNs and determine the appropriate priority of each path. The Multipath daemon presents the Pillar Axiom LUNs as virtual block devices in the Linux `device-mapper` framework. APM also provides the APM daemon, which monitors the state of the `device-mapper` and communicates with the Pillar Axiom software.

Note: You can use `device-mapper` and `multipath-tools` for many purposes other than managing multipathing for Pillar Axiom systems.

APM provides installation and start up scripts that automate several of the manual integration and configuration tasks that `device-mapper` usually requires. The automation includes:

- Bringing partitions on multipath devices online automatically
- Reordering startup scripts for optimum operation

Note: To determine how the Linux startup scripts have been altered, refer to the comments in the `/etc/init.d/axiompmd` file.

Supported Distributions

Pillar Axiom Path Manager (APM) is supported on Community Enterprise Operating System 6 (CentOS 6) platforms.

APM supports CentOS 6 distributions for the following architectures:

- x86-32 (32-bit x86 platforms)
- x86-64 (64-bit AMD and Intel platforms)

To determine the hardware platform on which your distribution is running, run the following command:

```
# uname -i
```

Compare the output of this command with the information in the following table:

Table 5 Supported hardware platforms

Hardware platform	Command output
x86	i386
AMD64/Intel EM64T	x86_64

To determine the Linux kernel installed on your system, run the following command:

```
# uname -r
```

Verify that the kernel identifier in the output of this command begins with the following value:

```
2.6.32-279
```

If the command output does not begin with the specified value, upgrade your operating system to the supported version before you install this release of APM. Alternatively, choose an APM that supports your operating system version.

APM Requirements

Pillar Axiom systems presenting LUNs to Community Enterprise Operating System 6 (CentOS 6) hosts using Pillar Axiom Path Manager (APM) must be running release 4.5 or higher of the Pillar Axiom software.

Operating Limits

Pillar Axiom Path Manager (APM) provides access over multiple data paths to LUNs defined on a Pillar Axiom system.

APM and the Pillar Axiom software limit the following aspects of this access.

Table 6 APM operating limits

APM capabilities	Maximum value
Target Pillar Axiom systems	Eight for each SAN host
Connect to SAN Slammer storage controllers	Four for each Pillar Axiom system
Connect to LUNs	256 for each Pillar Axiom system
Handle data paths	32 to each LUN
Handle FC HBA ports	32 for each SAN host

The Linux `device-mapper` and other parts of the Linux system limit the total number of devices and paths that can be used. Some of these limits can be tuned by the administrator. Please refer to the administrator's guide for your operating system for information on tuning your system limits.

Important! Not all combinations of the limits shown have been tested. Use care when operating a system that has been configured to run at or near these limits. The system may exhibit anomalies when all limits are exercised concurrently.

CHAPTER 2

Install APM

Prepare to Install APM

To ensure a successful installation of Pillar Axiom Path Manager (APM), perform the following tasks in sequence:

- 1 Read the APM Release Notes at the end of this document.
- 2 Ensure that the Pillar Axiom system is running release 4.5 or higher of the Pillar Axiom software.
- 3 If you are using FC or FC over Ethernet (FCoE) connections, verify that your FC SAN components and adapters are supported.
- 4 If you are using FC connections, confirm that your FC adapters are installed according to our instructions.

Supported adapters:

- Emulex
 - QLogic
 - Brocade
 - Xsigo
- 5 If you are using iSCSI connections, verify that your network is configured for iSCSI multipathing, and that you have configured the iSCSI software initiator correctly.
 - 6 If you are using iSCSI connections, confirm that any iSCSI-to-FC routers and IP switches you have installed are supported.
 - 7 Pre-configure the SAN host for Pillar Axiom integration.
 - 8 Verify installation of the `device-mapper`.

Related concepts

- [About Network Configuration for iSCSI Multipathing](#)
- [About Xsigo I/O Director Support](#)

Related references

- [Supported FC SAN Hardware Components](#)
- [Supported iSCSI-to-FC Routers](#)
- [Supported iSCSI Switches](#)

Related tasks

- [Install Emulex FC Adapters and Drivers](#)
- [Install QLogic FC Adapters and Drivers](#)
- [Install Brocade FC Adapters and Drivers](#)
- [Configure the iSCSI Initiator](#)
- [Start the iSCSI Software Initiator Service](#)
- [Verify Installation of the Device Mapper](#)
- [Pre-Configure the SAN Host for Pillar Axiom Integration](#)

Supported SAN Protocols

Pillar Axiom Path Manager (APM) supports FC and software-based iSCSI connections to the Pillar Axiom system.

Supported FC SAN Hardware Components

Pillar Axiom Path Manager (APM) supports a variety of FC HBAs, converged network adapters (CNAs), and SAN routers and switches.

Related references

- [Supported Emulex FC Adapters and Drivers](#)
- [Supported QLogic FC Adapters and Drivers](#)
- [Supported Brocade FC Adapters and Drivers](#)
- [Supported FC and FCoE Switches](#)

Supported FC Topologies

Pillar Axiom Path Manager (APM) supports the Point-to-Point (FC-P2P) and Switched Fabric (FC-SW) topologies. APM does not support Arbitrated Loop (FC-AL).

Maximum Queue Depth

The recommended maximum LUN queue depth for all SAN hosts attached to a Pillar Axiom system is 64. This value is the maximum number of outstanding I/O requests to each LUN on the Pillar Axiom system. Exceeding the maximum I/O queue depth can cause errors.

Supported Emulex FC Adapters and Drivers

Pillar Axiom Path Manager (APM) supports Emulex FC adapters drivers.

APM supports the following HBAs:

- LP10000
- LP10000DC
- LP10000ExDC-E
- LP11000
- LP11002
- LPe11000
- LPe11002
- LPe11004
- LPe12000
- LPe12002
- LPe12004

APM also supports the following Emulex FC over Ethernet (FCoE) converged network adapters (CNAs):

- OCe10102-F
- OCe11102-F

Note: These CNAs can also be used concurrently as Ethernet NICs.

The above HBAs and CNAs require the Emulex drivers that come with the Linux distribution.

Related references

- [Recommended Settings for Emulex FC Adapters](#)

Related tasks

- [Install Emulex FC Adapters and Drivers](#)

Recommended Settings for Emulex FC Adapters

We recommend the following settings for Emulex HBAs and converged network adapters (CNAs):

Table 7 Emulex adapter settings

Adapter setting	value
lpfc_lun_queue_depth	64
lpfc_devloss_tmo	60

Note: The installation instructions tell you how to set these parameters by editing the `/etc/modprobe.conf` file and rebooting the host. It is also possible to set these parameters using utilities supplied by Emulex. See the Emulex website for details.

Install Emulex FC Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your Emulex HBAs or converged network adapters (CNAs), be sure to follow the instructions in this section for installing the supported Emulex adapters and drivers.

- 1 Install Emulex HBAs and CNAs according to the instructions at the [Emulex support download page](http://www.emulex.com/downloads) (currently <http://www.emulex.com/downloads>).
- 2 Install the Emulex HBA and CNA drivers that come with the Linux distribution.
- 3 Edit the `/etc/modprobe.conf` file to set the options for the `lpfc` driver.

If an `options` line already exists for this driver, modify it to include the recommended settings. If an `options` line does not exist, add one in the following format:

```
options lpfc lpfc_devloss_tmo=60 lpfc_lun_queue_depth=64
```

- 4 Rebuild the kernel to ensure that any changes to the driver are picked up.

Use the following command:

```
# /sbin/new-kernel-pkg --mkinitrd --depmod --install
`uname -r`
```

- 5 Reboot the host.

Note: Emulex provides utilities that can be used to set the driver options without rebooting the host. See the Emulex website for details. If you use the Emulex utilities, you should still follow the steps above to update the configuration file and build a new kernel so that the options will be in effect after the next reboot.

Supported QLogic FC Adapters and Drivers

The Pillar Axiom Path Manager (APM) supports QLogic FC adapters and drivers.

APM supports the following QLogic HBAs:

- QLA2460
- QLA2462
- QLE2460
- QLE2462
- QLE2464
- QLE2560
- QLE2562
- QLE2564

APM also supports the following FC over Ethernet (FCoE) converged network adapters (CNAs):

- QLE8150
- QLE8152
- QLE8240
- QLE8242

Note: These CNAs can also be used concurrently as Ethernet NICs.

Important! QLogic has released a mandatory critical fix for the flash image on the QLE8240 and QLE8242 CNAs. See the QLogic website for details.

The above HBAs and CNAs require the QLogic drivers that come with the Linux distribution.

Related references

- [Recommended Settings for QLogic FC Adapters](#)

Related tasks

- [Install QLogic FC Adapters and Drivers](#)

Recommended Settings for QLogic FC Adapters

We recommend the following settings for QLogic HBAs and converged network adapters (CNAs):

Table 8 QLogic adapter settings

Adapter setting	Recommended value
ql2xloginretrycount	30

Table 8 QLogic adapter settings (continued)

Adapter setting	Recommended value
ql2xlogintimeout	60
qlport_down_retry	45
ql2xmaxqdepth	64

Note: The installation instructions tell you how to set these parameters by editing the `/etc/modprobe.conf` file and rebooting the host. It is also possible to set these parameters using utilities supplied by QLogic. Refer to the QLogic website for details.

Install QLogic FC Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your QLogic HBAs or converged network adapters (CNAs), be sure to follow the instructions in this section for installing the supported QLogic adapters and drivers.

- 1 Install QLogic HBAs and CNAs according to the instructions at the [QLogic support download page](http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/) (currently http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/).
- 2 Install the QLogic HBA and CNA drivers that come with the Linux distribution.
- 3 Edit the file `/etc/modprobe.conf` to set the options for the `qla2xxx` driver.

If an `options` line already exists for this driver, modify it to include the recommended settings. If an `options` line does not exist, add one in the following format:

```
options qla2xxx ql2xloginretrycount=30 ql2xlogintimeout=60
qlport_down_retry=45 ql2xmaxqdepth=64
```

- 4 Rebuild the kernel to ensure that any changes to the driver are included.

Use the following command:

```
# /sbin/new-kernel-pkg --mkinitrd --depmod --install
`uname -r`
```

- 5 Reboot the host.

Note: QLogic provides utilities that can be used to set the driver options without rebooting the host. See the QLogic website for details. If you use the QLogic utilities, you should still follow the steps above to update the configuration file and build a new kernel so that the options will be in force after the next reboot.

Supported Brocade FC Adapters and Drivers

Pillar Axiom Path Manager (APM) supports Brocade FC adapters and drivers.

APM supports the following Brocade HBAs:

- 415
- 425
- 815
- 825

APM also supports the following Brocade FC over Ethernet (FCoE) converged network adapters (CNAs):

- 1010
- 1020

Note: These CNAs can also be used concurrently as Ethernet NICs.

These HBAs and CNAs require the 3.1.0 (or later) driver, available from the [Brocade website](http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page) (<http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page>).

Download the packages and release notes from the Brocade website and follow Brocade's instructions to install the packages. Installing the packages will install the drivers.

Related references

- [Recommended Setting for Brocade FC Adapters](#)

Related tasks

- [Install Brocade FC Adapters and Drivers](#)

Recommended Setting for Brocade FC Adapters

We recommend the following setting for Brocade HBAs and converged network adapters (CNAs):

Table 9 Brocade adapter setting

Adapter setting	Recommended value
pathtov	60

Install Brocade FC Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your Brocade HBAs or converged network adapters (CNAs), be sure to follow the instructions in this section for installing the supported Brocade adapters and drivers.

- 1 Install Brocade HBAs and CNAs according to the instructions at the [Brocade support documentation download page](http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page#Docs) (currently <http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page#Docs>).
- 2 Download version 3.1.0 or later of the driver and utilities package from the [Brocade support driver download page](http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page#Drivers) (currently <http://www.brocade.com/services-support/drivers-downloads/adapters/Linux.page#Drivers>).
- 3 Follow the Brocade instructions to install the driver and utilities.
Be sure to update the adapter firmware if required.
- 4 Use one of the Brocade utilities to set the path time out value (`pathtov`) for each Brocade FC and FC over Ethernet (FCoE) port.

The path time out value should be set to 60. Refer to the Brocade documentation for details of how to set the value.

About Xsigo I/O Director Support

Xsigo I/O Director and associated host software provide proprietary FC emulation over InfiniBand.

The FC ports on the Xsigo I/O Director must connect to the FC fabric through an N_Port ID Virtualization (NPIV)-aware FC switch. Direct connection of the Xsigo I/O Director to the Pillar Axiom system is not supported.

In addition to FC emulation, Pillar Axiom Path Manager (APM) also supports the Ethernet emulation provided by the Xsigo solution. The Xsigo virtual Ethernet network interface controllers (NICs) are supported in the same ways as other Ethernet NICs.

Refer to the Xsigo documentation for instructions on setting up and configuring the Xsigo I/O Director and associated components. The InfiniBand host channel adapters (HCAs) and other equipment required for this solution are specified in the Xsigo documentation.

Related references

- [Supported Xsigo Drivers](#)

Supported Xsigo Host Drivers

Pillar Axiom Path Manager (APM) supports the Xsigo host drivers listed on the Xsigo Support Download Center page for [Linux Host Drivers](http://www.xsigo.com/supportdlc/category/linux-drivers/current-rhel6-x-drivers) (<http://www.xsigo.com/supportdlc/category/linux-drivers/current-rhel6-x-drivers>).

Refer to the [Xsigo Host Drivers Release Notes](http://www.xsigo.com/support/auth/_downloads/_files/docs/505-GA/Release_Notes_V505-LX_RevA_092612.pdf) (currently http://www.xsigo.com/support/auth/_downloads/_files/docs/505-GA/Release_Notes_V505-LX_RevA_092612.pdf) for more information.

Supported FC and FCoE Switches

For a list of supported FC and FC over Ethernet (FCoE) switches, choose one of the following options:

- Refer to the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Oracle Unified Storage Systems Documentation](http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html) website (<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>).
- Call Oracle Customer Support.

Note: Refer to the switch vendor's website for the most recent installation instructions, patches, and firmware.

Related references

- [Oracle Contacts](#)

Supported iSCSI Software and Hardware

Pillar Axiom Path Manager (APM) supports the following iSCSI software and hardware:

- The iSCSI software initiator included in the operating system distribution

- The iSCSI-to-FC routing features of the Cisco MDS 9000 family of routers and switches

Note: APM does not support iSCSI HBAs.

About Network Configuration for iSCSI Multipathing

You can reduce the impact of failures in your iSCSI network by ensuring that iSCSI connections are made through several NICs on the host.

We recommend using the normal IP configuration facilities to route connections to different ports on the Slammer storage controller through different host NICs. For example, if the host has two NICs available for iSCSI connections to a Slammer, you can set up routing to send all traffic to one port on each Slammer control unit (CU) through one NIC, and traffic to the other port on each Slammer CU through the other NIC.

The best way to ensure that your iSCSI paths use different physical connections depends on factors such as the number of available NICs in the host and the existing configuration of your IP network. We recommend using normal IP routing configuration techniques such as subnet configuration and explicit route configuration to distribute connections over multiple NICs. You can also use NIC bonding to provide redundancy on Ethernet connections between the host and the switch.

About Subnet Configuration

Groups of host NICs and Slammer ports can be configured into different subnets. Whenever the iSCSI initiator opens a connection to a Slammer port, that connection will automatically be routed through the host NIC in the same subnet as the Slammer port. For example, if a host has two NIC ports on the iSCSI network, one of the host NIC ports connected to one of the Slammer control unit (CU) ports could be placed in one subnet, and another host NIC port connected to the other Slammer CU port could be placed in a different subnet.

About NIC Bonding

Network interface card (NIC) bonding is another way to add redundancy to Ethernet networks.

With *NIC bonding*, also known as *channel bonding*, two or more NICs and their physical connections to the switches are logically bonded together and presented to the IP as a single virtual NIC. If one of the physical connections fails, the traffic

is transferred to another NIC without the IP layer or the layers above it knowing about the transfer.

This approach protects against low-level Ethernet failures, such as a faulty NIC or cable, between the host and its local IP switch. Because the redundancy is at a very low level in the protocol stack, the higher layers such as TCP/IP, iSCSI, and `device-mapper` are not aware that a transfer to a different NIC has taken place. To the IP protocol, the virtual bonded NIC appears as a single physical connection. To iSCSI, it appears as a single connection to each target port.

In short, the iSCSI, `device-mapper`, and APM layers are not aware of the physical path redundancy provided at the NIC bonding level and do not treat this redundancy as providing multiple paths to the Slammer storage controller. Multiple paths created by NIC bonding will not be reported as multiple paths by `device-mapper` or in the Pillar Axiom Storage Services Manager.

Configure the iSCSI Initiator

To support iSCSI connections, Pillar Axiom Path Manager (APM) requires the iSCSI software initiator that is included in your operating system distribution.

For complete installation and configuration instructions, refer to: `/usr/share/doc/iscsi-initiator-utils-*/README`. The steps below are an overview of the instructions in the `README` file.

Important! You must ensure that each iSCSI initiator on the iSCSI SAN has a unique name, and set this host's initiator name in the `/etc/iscsi/initiatorname.iscsi` file.

- 1 Set up IP addresses and network parameters on Slammer iSCSI ports.
This might include setting up Challenge Handshake Authentication Protocol (CHAP) on your system.
- 2 Configure iSCSI accordingly.
Example:
If you are using CHAP to communicate with the Slammer iSCSI ports, you may need to modify the `/etc/iscsi/iscsid.conf` file accordingly.
- 3 Start the iSCSI initiator.
Run the following command:

```
# service iscsi start
```
- 4 Discover the Slammer iSCSI target ports.
Examples:

- `# iscsiadm -m discovery -t sendtargets -p 172.20.191.11:3260`
- `# iscsi_discovery 10.0.0.22`

5 Log in to the Slammer iSCSI target ports.

Example:

```
# iscsiadm -m node --loginall=all
```

6 (Optional) Discover the other ports, depending on the desired design.

7 Restart the `iscsi` service to log in to the discovered target ports.

Run the following command:

```
# service iscsi restart
```

8 Use `chkconfig` to configure the `iscsi` service to start up at various run levels.

Related tasks

- [Start the iSCSI Software Initiator Service](#)

Start the iSCSI Software Initiator Service

After you have configured the iSCSI software initiator, set the iSCSI initiator to start at startup time and confirm that your iSCSI devices are visible.

1 Start the `iscsi` service.

Run the following command:

```
# service iscsi start
```

2 Configure the `iscsi` service to start at boot time.

Run the following commands:

```
# chkconfig --add iscsi
# chkconfig iscsi on
```

The first command checks that the scripts necessary to start and stop the service are available. The second command sets the service to start at the appropriate run levels.

3 Verify that the `iscsi` service is configured to start at boot time.

Run the following command:

```
# chkconfig --list iscsi
```

Result:

The following is a sample of the output of this command when the service is *not* configured to start at boot time:

```
iscsi 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

The following is a sample of the output of this command when the service *is* configured to start at boot time:

```
iscsi 0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

Supported iSCSI-to-FC Routers

iSCSI-to-FC routing features enable a host to use iSCSI to access LUNs on Pillar Axiom FC SAN Slammers.

Pillar Axiom Path Manager (APM) supports the iSCSI-to-FC routing features of the Cisco MDS 9000 family of multilayer directors and fabric switches. The only supported iSCSI-to-FC routing solution is the solution provided by this family of switches.

The iSCSI-to-FC features were tested on Cisco MDS SAN-OS Release 3.0 (2a).

For more information on these features, refer to the [Cisco documentation](http://www.cisco.com/) (http://www.cisco.com/).

Configure the Cisco MDS switch as an iSCSI-to-FC router to use with APM and the Pillar Axiom system.

Related tasks

- [Configure the iSCSI-to-FC Router](#)

Supported iSCSI Switches

For a list of supported iSCSI switches:

- Refer to the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Oracle Unified Storage Systems Documentation](http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html) website (http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html).
- Call Oracle Customer Support.

Note: Refer to the vendor's website for the most recent installation instructions, patches, and firmware.

Related references

- [Oracle Contacts](#)

Configure the iSCSI-to-FC Router

The iSCSI-to-FC routing features of the Cisco MDS 9000 family of multi-layer directors and fabric switches require configuration to work with the Pillar Axiom Path Manager (APM) and the Pillar Axiom system.

For more information on these features, refer to the [Cisco documentation](http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/3_0/fmcfg/index.htm) (http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/3_0/fmcfg/index.htm).

- 1 Present the ports on the Slammer storage controller as iSCSI targets.

Choose **Dynamic Mapping** or **Static Mapping**. However, we recommend that you use dynamic mapping because the main Cisco features for static mapping requirements are supplied by APM and the Pillar Axiom Storage Services Manager.

- 2 Present the iSCSI hosts as virtual FC hosts.

The hosts must be presented in transparent initiator mode (not in proxy-initiator mode). When you assign World Wide Names (WWNs) for the iSCSI initiators, use the static mapping mechanism.

After you configure the switch, APM on the iSCSI hosts interacts with the Pillar Axiom systems in exactly the same way as when both hosts and Slammers use the same SAN protocol.

Pre-Configure the SAN Host for Pillar Axiom Integration

Before you install the Pillar Axiom Path Manager (APM) software, prepare your SAN host connections with the Pillar Axiom system.

Prerequisites:

Verify that your system has the following characteristics:

- At least one SAN Slammer storage controller that has FC or iSCSI ports
- Supported HBA and converged network adapter (CNA) drivers

- Ethernet connections to the management ports on the Pilot management controller
- A network configuration that allows an application on the SAN host to connect to TCP port 26004 on the Pilot

Prepare your SAN host components so that they will be able to connect with the Pillar Axiom system.

- 1 Verify that all FC and iSCSI components and software are installed on the SAN host according to the instructions in this guide.
- 2 Set up the physical connectivity and any required switch zoning for the SAN.

Proper setup is needed so all required host ports can access the Slammer ports on the Pillar Axiom system.

- 3 If you are using iSCSI connections, verify that your network is configured for iSCSI multipathing, and that you have configured the iSCSI software initiator correctly.
- 4 Launch the Pillar Axiom Storage Services Manager and navigate to the **Storage > SAN > Hosts** page.
- 5 Verify the SAN connections.

Each host port should display individually as follows:

Table 10 Characteristics of SAN connections to host ports

Column label	Port type	
	FC	iSCSI
Host Name:	Hostname Unknown	IQN (iSCSI Qualified Name)
Host Port:	WWN	IP address

Related concepts

- [About Network Configuration for iSCSI Multipathing](#)

Related references

- [Related Documentation](#)

Related tasks

- [Configure the iSCSI Initiator](#)
- [Start the iSCSI Software Initiator Service](#)
- [Configure SAN Host Access to the Pillar Axiom LUNs](#)

Verify Installation of the Device Mapper

The Linux 2.6 `device-mapper`, among other uses, maps a logical block device on the SAN host to a Pillar Axiom LUN. The Pillar Axiom Path Manager (APM) daemon monitors the `device-mapper` and uses it to provide multipathing.

APM requires the `device-mapper` and `multipath-tools` version that comes with the Linux distribution.

- 1 Verify that `device-mapper` is installed.

Run the following command:

```
# rpm -qa | grep device-mapper
```

Result:

The versions of all installed packages that include `device-mapper` in their names are listed.

- 2 If the packages mentioned above are not installed, install them.

You can install the packages from your Linux installation CDs or the operating system vendor website.

About APM Software Installation

After you prepare your SAN for Pillar Axiom Path Manager (APM) as described in the first half of this chapter, you can download and install the APM software. After the installation, configure your SAN host and the Pillar Axiom software so that they can work together.

The APM installation requires that you download the APM software installation package from the Oracle Technical Network (OTN) and install the software on your system.

After you install APM, configure access from the SAN host to Pillar Axiom LUNs.

Related tasks

- [Download the APM Software](#)
- [Install \(or Update\) the APM Software](#)
- [Remove the APM Software \(Optional\)](#)
- [Configure SAN Host Access to the Pillar Axiom LUNs](#)

Download the APM Software

The software and documentation are distributed as a single zip archive that is available for download from the Oracle Technical Network.

Prerequisite: Join the [Oracle Technology Network](http://www.oracle.com/technetwork/community/join/why-join/index.html) (<http://www.oracle.com/technetwork/community/join/why-join/index.html>) to gain access to software and documentation downloads.

- 1 On the Oracle Technology Network, open the [Pillar Axiom Downloads](http://www.oracle.com/technetwork/server-storage/san-storage/downloads/index.html) page (<http://www.oracle.com/technetwork/server-storage/san-storage/downloads/index.html>) in your browser.
- 2 Under the Pillar Axiom Downloads title, select **Accept License Agreement** to be able to download the software package.
- 3 Click the name of the software package to download.
- 4 Extract the contents of the software bundle archive to your SAN host.

The archive contains software installation packages for all supported hardware platforms, as well as documentation, for the specified APM. Extract the package for your hardware platform and the documentation.

After you download the software, you can install it on your host system.

Install (or Upgrade) the APM Software

Install the software that you extracted to your SAN host so you can use Pillar Axiom Path Manager (APM).

Prerequisite: Before you install the APM software, verify that your system meets all hardware and software prerequisites mentioned earlier in this document.

Two steps in the installation process enable you to keep your original multipath configuration file (`/etc/multipath.conf`). You must complete these two steps to continue using your specific configurations for managing devices other than APM-managed Pillar Axiom LUNs.

- Because the configuration file will be deleted, Step 1 asks you to save a copy of that file before you begin the APM installation or update task.
- Step 4 asks you to merge any changes you made to the original configuration file into the new file.

Tip: We strongly recommend that you follow the above mentioned two steps when you have previously added or modified entries in the `/etc/multipath.conf` file.

- 1 If you previously configured `multipath-tools` on your system, save a copy of the current `/etc/multipath.conf` file in a different directory.
Saving a copy allows you to merge your customizations into the new default file that will be created.
- 2 Remove any previous versions of APM on your system.

Use the following command:

```
# rpm -e axiompmp
```

- 3 Install the APM package.

Use the following command:

```
# rpm -ivh axiompmp-version.rpm
```

Note: In the command above, *version* is the name of the release version you downloaded.

APM files and scripts are installed in the `/opt/pillar` directory.

- 4 After APM and the new `/etc/multipath.conf` is installed, merge in any previous configuration data for `device-mapper` from the copy of the configuration file you saved in Step 1.

Merging those changes allows you to retain previous configuration settings to manage devices other than APM.

- 5 Restart the host.

Configure SAN Host Access to the Pillar Axiom LUNs

The LUN configuration procedure provides the SAN host with access to Pillar Axiom LUNs. Refer to the *Pillar Axiom Administrator's Guide* for instructions on how to create LUNs.

Prerequisites:

- Pre-configure the SAN host for Pillar Axiom integration.
- Install the Pillar Axiom Path Manager (APM) software.

Use the Pillar Axiom Storage Services Manager to verify that your SAN host has access to the Pillar Axiom LUNs, and to configure any additional LUNs that are required.

- 1 In the Pillar Axiom Storage Services Manager, navigate to the **Storage > SAN > Hosts** page.
- 2 Verify that the individual entries for the host ports have been replaced with a single entry under the host name.

Examples:

Figure 2 Example host ports before APM installation

Host Name	Host Port	Type	AxiomONE Path Manager	Number of LUNs	Host Port Status
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:36:84:6e	FC	Not Registered	0	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:36:84:6f	FC	Not Registered	0	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:36:85:20	FC	Not Registered	6	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:41:32:c3	FC	Not Registered	0	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:41:32:c4	FC	Not Registered	0	Connected
<input type="checkbox"/> iqn.1987-05.com.cisco:01.eca9a9b8d555	192.168.2.93	iSCSI	Not Registered	0	Connected
	192.168.2.94	iSCSI			Connected

Figure 3 Example host ports after APM installation

Host Name	Host Port	Type	AxiomONE Path Manager	Number of LUNs	Host Port Status
<input type="checkbox"/> happy	10:00:00:00:c9:36:84:6e	FC	Communicating	0	Connected
	10:00:00:00:c9:36:84:6f	FC			Connected
	192.168.2.93	iSCSI			Connected
	192.168.2.94	iSCSI			Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:36:85:20	FC	Not Registered	6	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:41:32:c3	FC	Not Registered	0	Connected
<input type="checkbox"/> Hostname Unknown	10:00:00:00:c9:41:32:c4	FC	Not Registered	0	Connected

Note: The Hosts page may display differently in your version of Pillar Axiom Storage Services Manager.

You will see one or more of the following **Pillar Axiom Path Manager Status** and **Host Port Status** messages on the Hosts page:

APM Status **Communicating:** The host control path is currently logged into the Pilot.

Note: **Communicating** status is required for the APM control path to report path status, configure load balancing, and use the Pillar Axiom system to collect APM diagnostic logs.

Not Registered: A control path from an APM host with this name has never logged into the Pilot.

Not Communicating: The APM host control path has previously logged into the Pilot, but it is not currently logged in.

Host Port Status **Connected:** The host SAN connection is logged in to the SAN Slammer.

Not connected: The host SAN connection is not logged in to the SAN Slammer.

See the Pillar Axiom Storage Services Manager Help for information about the remaining fields on the Hosts page.

- 3 As needed, create new LUNs on the Pillar Axiom system for the SAN hosts.
- 4 As needed, set up mappings of the LUNs to the new host entries.
- 5 Make any changes visible to the SAN host by following the recommendations for SAN dynamic reconfiguration.
- 6 (Optional) List the multipath devices.

Run the following commands consecutively:

```
# /sbin/multipath -v3
# /sbin/multipath -ll
```

The first command (`multipath -v3`) populates the path information, and the second command (`multipath -ll`, lower-case letters L) lists the state of the paths.

- 7 In the Pillar Axiom Storage Services Manager, navigate to the **Storage > SAN > Hosts** page.
- 8 Click the name of the new host and, on the Host Information page, verify the APM software version.

- 9 Click the **LUN Connections** tab and verify that the host and LUN connections are as expected.

The column titled **LUN Name on Host** should show the Linux disk names that APM allocates to each LUN.

Important! With a few exceptions (such as when you partition and format the LUN on the SAN host), you should use only these Linux disk names to access and configure the LUNs on the host. Linux creates device names for each individual path to a LUN, but almost all configuration and administration tasks should be done using the `/dev/mapper` name shown in the Pillar Axiom Storage Services Manager.

If you plan to use a LUN disk, you must partition and format the LUN disk on the SAN host.

Related references

- [\[13746418\] SAN Dynamic Reconfiguration](#)
- [Related Documentation](#)

Related tasks

- [Download the APM Software](#)
- [Install \(or Update\) the APM Software](#)
- [Remove the APM Software \(Optional\)](#)
- [Pre-Configure the SAN Host for Pillar Axiom Integration](#)
- [Install \(or Update\) the APM Software](#)
- [Partition and Format the LUN Disk \(Optional\)](#)

Partition and Format the LUN Disk (Optional)

Follow these instructions if you plan to use the Linux `fdisk` or `parted` utilities to partition and format the disk on which a LUN resides.

The `fdisk` utility cannot be used with devices listed in the `/dev/mapper` directory. Instead, use `fdisk` on one of the underlying paths, and then run the scripts to restart the `multipath-tools` and APM daemons to notify `device-mapper` that a `/dev/mapper` device has a new partition.

- 1 Identify one of the highest priority paths to the LUN using the output of the `multipath -ll` command.

Example:

In the output below, you could identify either the `/dev/sdd` path or the `/dev/sdt` path:

```
2000b08008e001305 dm-0 Pillar,Axiom 600
size=50G features='1 queue_if_no_path' hwhandler='0' wp=rw
| -+- policy='round-robin 0' prio=4000000 status=active
|   |- 4:0:1:1 sdd 8:48 active ready running
|   `-- 5:0:1:1 sdt 65:48 active ready running
| -+- policy='round-robin 0' prio=3000000 status=enabled
|   |- 8:0:0:1 sdj 8:144 active ready running
|   `-- 6:0:0:1 sdk 8:160 active ready running
| -+- policy='round-robin 0' prio=2000000 status=enabled
|   |- 4:0:0:1 sdb 8:16 active ready running
|   `-- 5:0:0:1 sdf 8:80 active ready running
| `+- policy='round-robin 0' prio=1000000 status=enabled
|   |- 9:0:0:1 sdl 8:176 active ready running
|   `-- 7:0:0:1 sdp 8:240 active ready running
```

- 2 Use `fdisk` to partition one of the highest priority paths identified in Step 1.
- 3 Run the following command to restart the `device-mapper`:

```
# /etc/init.d/multipathd restart
```

- 4 Run the following command to restart the APM daemon:

```
# /etc/init.d/axiomcmd restart
```

- 5 Verify that the new partitions are listed as:

```
/dev/mapper/LUNpx
```

Where *LUN* is the LUN identifier, and *x* is the partition number used in Step 2.

Note: The sequence *p* appears between the LUN identifier and the partition number.

Example:

In the example above, if only partition 1 were created with `fdisk`, it would appear as follows:

```
localhost$ ls -l /dev/mapper/2000b08008e001305*
brw-rw---- 1 root disk 253,  2 Jul  7 12:02
2000b08008e001305
brw-rw---- 1 root disk 253,  3 Jul  7 15:12
2000b08008e001305p1
```

6 Format the new partition.

Example:

To format the partition created in the example above, you might run the following command:

```
# mke2fs -j /dev/mapper/2000b08008e001305p1
```

7 Mount the filesystem.

For detailed information, refer to your Linux documentation.

About Load Balancing Configurations

Pillar Axiom Path Manager (APM) can be configured for static or round-robin load balancing. You can configure load balancing separately for each LUN.

In static load balancing, the software selects the best available path and all commands are sent over that path until the path is no longer operational or a better path becomes available. Then, a failover to another appropriate path is initiated.

In round-robin load balancing, commands are sent by turn over the best available paths. This ensures that LUN commands are evenly distributed over any path that is available to access the LUNs.

Load balancing allows the paths to share the load in different ways:

- Balances access to a LUN across all optimized Slammer ports available for that LUN
- Balances access from a host across the host's HBA channels

To configure round-robin or static load balancing through the Pillar Axiom Storage Services Manager, refer to the *Pillar Axiom Administrator's Guide*.

Related references

- [Related Documentation](#)

About Path Selection

Pillar Axiom Path Manager (APM) selects the best paths to access Pillar Axiom LUNs.

Path selection is based on three factors:

- Path optimization state
- Path performance
- Path availability

Path optimization is determined by the Slammer control unit (CU) that the path uses. An *optimized path* is a path that connects through the Slammer CU on which the LUN is currently resident, or *homed*. A *non-optimized path* is a path that connects through the alternate Slammer CU. The LUN can be accessed through the alternate CU, but the LUN is not resident on the alternate CU. Optimized paths are always preferred, but if an optimized path is not available, the non-optimized path can be used temporarily.

Path performance is determined by how quickly and reliably a path can transfer I/O traffic to and from a LUN. Generally, FC paths perform better than iSCSI paths, so FC paths are preferred over iSCSI paths.

Path availability is determined by the ability of the path to transfer I/O traffic. An *available* path is fully functional. If the path stops working, it is considered *unavailable*.

These factors determine how the paths to each Pillar Axiom LUN are divided into groups.

APM groups the paths in the following order of preference:

- First, FC optimized
- Next, iSCSI optimized
- Next, FC non-optimized
- Finally, iSCSI non-optimized

For each LUN, the currently configured load balancing algorithm is used to select paths from the most preferred group that has paths available. Only paths from a single group are used at any one time.

When an active path fails, I/O traffic is transferred to a different path. I/O performance is reduced for a short time while the operating system recognizes

the failure and makes the path transfer. After the path transfer is complete, I/O performance improves.

If the failing path is optimized and the new path is non-optimized, I/O performance might continue to be reduced after path transfer because a lower-performance path is in use. Within a few minutes of traffic being transferred to a non-optimized path, the Pillar Axiom system attempts to reconfigure the LUNs to use an optimized path. After transfer to an optimized path succeeds, I/O performance improves.

Remove the APM Software (Optional)

When you uninstall Pillar Axiom Path Manager (APM), support for multiple paths to Pillar Axiom LUNs is removed.

- 1 Run the following commands to stop and remove the APM daemon:

```
# /etc/init.d/axiompmd stop
# rpm -e axiompmd
```

After the APM software (`axiompmd`) has been removed, the host appears not to be communicating in the Pillar Axiom Storage Services Manager, and the multipath behavior stops working.

Important! If you want to continue LUN access without multipathing, we recommend that you reconfigure the SAN so that only a single path exists from the host to each Pillar Axiom SAN Slammer storage controller. You may need to reconfigure LUN mapping on the Pillar Axiom system to map the LUNs to the host port or iSCSI Qualified Name (IQN). You also may need to reconfigure the host to access the LUNs through new device names.

CHAPTER 3

APM Release Notes

New in this Release

- Pillar Axiom Path Manager (APM) 3.3 supports Community Enterprise Operating System 6.3 (CentOS 6.3) in this release.
- Xsigo virtual FC and Ethernet adapters are supported in this release.

Known APM Issues

There are no known Pillar Axiom Path Manager (APM) issues for this release.

Known Pillar Axiom Issues

The following issues might be associated with the version of the Pillar Axiom software you are using.

Table 11 Known Pillar Axiom issues

Issue	Workaround or planned fix
<p>[13759030] If an iSCSI initiator is connected to a port on a Slammer, and that Slammer CU is powered off, the LUN and Host GUI pages continue to show the connection status for the iSCSI initiator as <i>Connected</i>.</p>	<p>This issue is fixed in release 5.0 of the Pillar Axiom software.</p>
<p>[13759805] If more than 256 SAN LUNs are configured on a Pillar Axiom system, the Pillar Axiom Storage Services Manager may send invalid messages to the APM daemon running on SAN host systems.</p> <p>When this happens, the control path connection between APM and the Pillar Axiom will continually move between <i>Communicating</i> and <i>Not Communicating</i> states. This prevents features that depend on the APM control path (such as setting the load balancing algorithm) from working properly. The data path, which manages LUN access from the host, is not affected.</p>	<p>This issue is fixed in release 5.0 of the Pillar Axiom software.</p>
<p>[13764561] The Pillar Axiom Storage Services Manager (GUI) sometimes continues to display host paths that are no longer valid after APM stops.</p> <p>This is because APM is no longer communicating path information to the GUI, so the GUI continues to display the host paths as they were while APM was running.</p>	<p>None.</p> <p>This issue will be fixed in a future release of the Pillar Axiom software.</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
<p>[13764609, 13762326] While the system is recovering from temporary use of non-optimized paths to the alternate CU on a Slammer, you may notice a decrease in I/O performance between a LUN and a Slammer CU.</p>	<p>This issue is fixed in release 4.5.1 and 5.2.1 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.5.1 or 5.2.1, you can take the following actions:</p> <ol style="list-style-type: none"> 1 Follow the instructions in the <i>Pillar Axiom Administrator's Guide</i> to rehome the LUN to the alternate CU on that Slammer. 2 Rehome the LUN again to the original CU on that Slammer. <p>Contact Oracle Customer Support for assistance.</p>
<p>[14576187] When you remove APM from a host and then delete the host entry in the Pillar Axiom Storage Services Manager (GUI), that host's initiators are considered unassociated. If you then attempt to use the GUI to associate one of the host initiators to another host, the association fails with the error <code>Unsatisfied request due to internal error</code>.</p>	<p>Use the Pillar Axiom CLI, not the GUI, to form the association.</p> <p>This issue will be fixed in a future release of the Pillar Axiom software.</p>
<p>[14581579] A defect in the Pillar Axiom software prevents proper creation of an APM host entry in the management server when there are more than 1024 LUN mappings to the initiators associated with the host, and the mappings differ between the initiators. You can recognize this defect in the Pillar Axiom Storage Services Manager (GUI) when the following symptoms occur together:</p> <ul style="list-style-type: none"> • The APM host is not communicating. 	<p>Take the following actions:</p> <ol style="list-style-type: none"> 1 Stop the APM service on the host. 2 Manually reconcile the mappings on all of the host initiators so that all initiators see the exact same set of LUNs via the same LUN number. 3 Restart the APM service to enable the host to configure itself. <p>This issue will be fixed in a future release of the Pillar Axiom software.</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
<ul style="list-style-type: none">The system displays the following TaskFailed system alert: <code>UNSATISFIED_REQUEST_PMI_COMMUNICATION_ERROR</code>	
<p>[14747038] The Pillar Axiom software is sometimes unable to configure the load balancing setting for LUNs on APM hosts. Attempts to change the load balancing setting through the Pillar Axiom Storage Services Manager (GUI) or CLI fail with the following error:</p> <pre>Unsatisfied request due to internal error</pre>	<p>None.</p> <p>This issue will be fixed in a future release of the Pillar Axiom software.</p>

Known Operating System Issues

The following operating system issues can have an impact on running Pillar Axiom Path Manager (APM) on Linux systems.

Filesystems Marked Read-Only When Access to Boot LUN Is Lost

When a SAN host loses access to all paths to a boot LUN, Linux marks the filesystem on the client as read-only.

Linux is sensitive to losing all paths to a boot LUN. When a SAN host loses access to all paths to a boot LUN, to prevent data corruption, the Linux system marks the filesystem on the client as read-only. This behavior is as designed and would occur regardless of whether the Pillar Axiom Path Manager (APM) is installed on the SAN host. To improve the path recovery time, we recommend that you take the following actions:

- Modify the `/etc/multipath.conf` file and set `failback immediate`.
- Configure the host to minimize the chances of all paths being lost at the same time.

Access to LUNs

In most cases, you should only access LUNs through `device-mapper` or individual partition paths.

The Linux `device-mapper` creates paths of the form `/dev/mapper/2000b08003d001321` to represent multipath Pillar Axiom LUNs. Individual partitions on a LUN have names of the form `/dev/mapper/2000b08003d001321p3`. With the exception of the situation described in the instructions for partitioning and formatting LUN disks, you should access multipath LUNs and their partitions exclusively through these paths.

If a multipath partition is configured into `/etc/fstab`, omit the sixth field of its entry (`fs_passno`), or set it to 0 to prevent `fsck` from running automatically on the partition during a system boot. This is because `device-mapper` is not yet configured at the time `fsck` runs during boot, so the multipath devices are not accessible.

Important! Failure to disable `fs_passno` will cause host boot failure.

Related tasks

- [Partition and Format the LUN Disk \(Optional\)](#)

[13746418] SAN Dynamic Reconfiguration

Linux does not automatically detect storage LUNs after a dynamic reconfiguration. Dynamic reconfiguration is the addition, deletion, growing, resizing, or cloning of one or more LUNs attached to a host.

Follow the instructions for dynamic reconfiguration in the Online Storage Management chapter of the [Storage Administration Guide](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch24.html) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch24.html). If you continue to have problems with dynamic reconfiguration, the following steps will help you handle exceptions.

Linux requires a series of steps, including a potential host reboot, when a LUN is dynamically reconfigured. After Linux detects the new LUNs, you must restart Pillar Axiom Path Manager (APM) to update the Pillar Axiom system with the new LUN status.

LUNs Added Dynamically

In most systems, a newly added LUN is immediately visible on the host without a rescan. However, due to inconsistent device driver behavior on some hosts, if the added LUN is not visible, a rescan usually makes it visible. A rescan normally involves an I/O reset.

Create a new LUN on the Pillar Axiom system and assign it to the host. If running the rescan script does not bring up the LUNs, you may need to assign a new LUN number that the Linux operating system does not incorrectly consider already in use. To correct this situation, modify the host LUN number in the Pillar Axiom Storage Services Manager. Assign a new, unique value that falls within the range of permitted values. If necessary, rescan to add the LUN.

Both Emulex and QLogic provide rescan scripts that may help in dynamically configuring LUNs. See your adapter vendor's website for information on any tools they provide for scanning the SAN.

LUNs Deleted Dynamically

Deleting a LUN prevents the LUN from being visible from the host. This includes deleting LUN mapping and LUN masking. In general, LUN deletion disrupts normal function of the Linux multipath framework and must be planned.

If a LUN is deleted, it may appear as either a 20000000000000 entry or as the original LUID with `Path down` messages. These entries may persist until the host is rebooted.

To avoid disruption, you may blacklist the LUN. Refer to your Linux documentation.

The host usually picks up the deleted LUN, and it is deleted from the `/dev/mapper` table. However, this may not occur on all platforms consistently. If you want to view the device-mapper LUN mapping table, start the `multipathd` shell by running the following command:

```
# /sbin/multipathd -k
```

To delete a LUN, we recommend shutting down the host, deleting the LUN or LUN mapping from the Pillar Axiom system, and then restarting the host. If this procedure is not possible, you may want to run the following procedure.

Important! The following procedure will interrupt I/O and may require an immediate reboot of your host. In some cases, this may require a power cycle of the host to recover.

- 1 Copy the following and run it as a script:

```
#!/bin/bash
# Must be run as root
/etc/init.d/axiompmd stop
/sbin/multipath -F
/etc/init.d/multipathd stop
# RESCAN SCRIPT FROM QLOGIC / EMULEX
# Please modify the following line based on your rescan
script location
/usr/bin/ql-dynamic-tgt-lun-disc.sh -s -r
/etc/init.d/multipathd start
/etc/init.d/axiompmd start
/sbin/multipath -v3 -ll
```

Tip: The rescan script might require your interaction.

- 2 Be prepared to reboot the host as soon as possible after deleting LUNs in case something goes wrong.

If a LUN that is visible to a Linux 2.6 host is deleted from the Pillar Axiom system, and the `/sbin/multipath -F` or `/sbin/multipath -f` command is run before rebooting the host, the device-mapper configuration map may become unusable and all access to LUNs may be lost due to a bug in the Linux device-mapper code. If this occurs, the only way to recover is to reboot the host. Rebooting the host will flush any deleted path entries from the device table.

LUNs Resized Dynamically

When you resize a LUN, a host reboot is necessary due to the constraints in the Linux device-mapper. You must stop the iSCSI services before you reboot the host.

However, if you follow the procedures documented in the Online Storage Management chapter of the [Storage Administration Guide](http://) (<http://>

docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch24.html) to force your device drivers to recognize the resized LUN, the device-mapper may recognize the resized LUN without a host reboot. The ability to use the resized LUN is a function of the host filesystem.

Clone LUNs Added or Deleted Dynamically

The procedures for adding or deleting LUNs described above also apply for Clone LUNs.

[13758951] Multipath Failed Path Errors

The `multipath` command may return errors that indicate only that there are failed paths.

If paths are in a failed state, the `multipath` command returns the following error messages:

```
multipath -ll 2000b08005c001259
9:0:0:3: sg_io failed status 0x8 0x1 0x0 0x0
9:0:0:3: Unable to get INQUIRY vpd 1 page 0x0.
error calling out /sbin/scsi_id -g -u -s /block/sdaf
8:0:0:3: sg_io failed status 0x8 0x1 0x0 0x0
```

These error messages indicate only that there are failed paths in the multipath device map. The multipath device map shows paths that are failed and active, and `sg_io failed` refers to the fact that SCSI generic (`sg`) devices do not exist for iSCSI device paths that are failed. These errors indicate that the system is responding correctly.

You must fix the failed paths or, if these errors occur during failover testing, recognize that this is normal and expected Linux multipath behavior.

Note: The `multipath` command can be invoked automatically by the system at various times, so it is possible for messages like these to be seen whenever paths are in an unusual state, such as during dynamic reconfiguration.

[13764910] Non-Optimized Access Events

Non-optimized access (NOA) events may be logged on the Pillar Axiom system for up to ten minutes after a Slammer control unit (CU) failover and failback, after all paths to a CU fail and subsequently recover, or after a new LUN is mapped to the host.

After a CU failover and failback or a new LUN is mapped to the host, it can take some time for Linux to recover all paths and bring them back into use. During this period the host may continue to access the LUNs through paths which did not fail. This may cause short bursts of NOA to some LUNs. Once Linux brings the paths fully back into use, the host will synchronize again with the Slammer on the best paths to use, and NOA events will no longer be logged.

[13761123] LUN State Not Updated

Due to a Linux defect in the `device-mapper` packages, `multipath-tools` may not accurately update the state of the LUN.

When the LUN state is not accurately updated, the Pillar Axiom system may not reflect the path count properly during a port enable or port disable test.

To correct this problem, restart the `device-mapper` to update the path status.

This is a Linux bug that the operating system vendor will fix in future updates.

[13761100] Too Many Files Open

The `multipath` command will issue an error indicating that too many files are open when the number of open files exceeds the configured host system limit.

To determine the number of files open, run the following command:

```
# lsof | wc -l
```

Compare the result of that command with the host system limit. To determine the host system limit, run the following command:

```
# ulimit
```

If necessary, increase the system limit by following the instructions in the operating system documentation. In addition, set the `max_fds` value in the `/etc/multipath.conf` file to an appropriate limit.

A system reboot is needed to ensure that the system limits are updated.

[13762235] Too Many Paths Cause Device Mapper to Hang

The `device-mapper` daemon occasionally hangs if operating system limits on the number of paths or devices, or other related limits, are exceeded.

A few paths appear to be missing in the host, and the `multipathd show maps topology` command hangs.

It may be necessary to restart the multipath daemon to correct the problem. Use the following command:

```
# service multipathd restart
```

Some of these limits can be tuned by the administrator. Please refer to the administrator's guide for your operating system for information on tuning your system limits.

[13762937, 13766418] LUNs Not Shown

Due to a Linux `device-mapper` defect, some LUNs are not shown in the output of the `/sbin/multipathd show maps topology` command.

The unused LUNs exist and are otherwise usable. This only affects reporting of LUNs to the Pillar Axiom system and does not affect I/O or usage of the device. The Pillar Axiom Path Manager (APM) is unable to update the host path status on the Pillar Axiom system, but I/O continues as expected. As a result, APM displays incorrect information.

[13764931, 13792540] Device Mapper Driver Not Loaded

The Linux operating system does not automatically load the `device-mapper` driver by default upon restarting the `multipathd` service, so it may be necessary to load the driver manually before installing Pillar Axiom Path Manager (APM).

The multipath driver must always be loaded before APM is installed.

If APM is installed prior to loading the `device-mapper` driver, the following error message appears:

```
# DM multipath kernel driver not loaded
```

The `device-mapper` driver can be loaded in one of two ways:

- `mpathconf --with_module y`

Refer to the `mpathconf` man pages for an explanation.

- Manually load the `device-mapper` using the following command:

```
# modprobe dm_multipath
```

The first method is preferred.

[13819519] CNAs Require Firmware Update

Before you install the drivers for any Converged Network Adapters (CNAs), be sure to update the firmware. Refer to the installation instructions on the vendor website.

[13959093] LUNs Not Seen

The Pillar Axiom system might not be able to see all LUNs reported from the host through the APM daemon.

After certain failover situations, the `multipathd show maps topology` command on the host does not list all the storage LUNs and paths.

This is a known Linux problem that affects only the information display, not the actual I/O of the LUNs.

[14708044] Xsigo Driver Fails to Restore Paths

After failover and failback of a host using a Xsigo virtual FC adapter, traffic does not come back to its original paths when the failed CU is back online.

Workaround: To recover the paths, run a Xsigo host rescan, and then run the `multipath -v3` command. This restores the paths and I/O returns to the optimized paths.

[14708089, 14711677] Adding or Deleting LUNs Requires Reboot

Addition or deletion of LUNs on a Xsigo virtual HBA (vHBA) or a Xsigo device connected to a Linux 6.3 host requires a host reboot. As a result, we recommend that you add or delete LUNs only when a host reboot can be accommodated.

Resolved APM Issues

There are no resolved APM issues for this release.

Additional Notes

Mount iSCSI Filesystems

We recommend that iSCSI filesystems be auto-mounted with the `_netdev` option. If the iSCSI filesystem is manually mounted, it must be unmounted manually before a reboot to avoid hangs.

Filesystems installed on iSCSI devices cannot be automatically mounted at system reboot because the IP network is not yet configured at mount time. However, the driver provides a method to auto-mount these filesystems as soon as the iSCSI devices become available, after the IP network is configured.

To auto-mount a filesystem installed on an iSCSI device:

- 1 List the iSCSI partitions to be automatically mounted in `/etc/fstab`.
- 2 For each filesystem on each iSCSI device, enter the logical volume on which the filesystem resides.

The mount points must exist for the filesystems to be mounted. For example, the following `/etc/fstab` entry will mount the iSCSI devices with partition `p1` specified by the LUID `20000121390p1`:

```
#device mount FS mount backup fsck
#to mount point type options frequency pass
LABEL=/ / ext3 defaults 1 1
LABEL=/boot /boot ext3 defaults 1 2
/dev/mapper/20000121390p1 ext2 _netdev 0 0
```

- 3 Restart the system.

Result:

The `netfs` initialization script mounts all filesystems with the `_netdev` mount option. Run this initialization script after the networking and iSCSI scripts to ensure that the system is ready to mount these devices.

Due to variable network delays, targets may not always become available in the same order from one boot to the next. Thus, the order in which iSCSI devices are mounted may vary and may not match the order the devices are listed in `/etc/fstab`. You should not assume mounts of iSCSI devices will occur in any particular order.

Index

A

- additional notes *63*
- additional resources *8*
- APM
 - how to
 - download software *38*
 - remove software *49*
 - update software *39*
 - interaction with Pillar Axiom *14*
 - pre-installation requirements *22*
- APM software
 - control path *16*
 - data path *17*
 - definition *10*
 - features *10*
 - known issues *51*
 - operating limits *21*

B

- Brocade CNAs
 - install *29*
 - settings *28*
 - supported *28*
- Brocade HBAs
 - install *29*
 - settings *28*
 - supported *28*

C

- Cisco MDS switch
 - configure *35*
- CNAs
 - Brocade *28*
 - QLogic *26*
- configuration
 - iSCSI multipathing *31*
 - iSCSI-to-FC router *35*
 - Pillar Axiom LUNs *41*
 - SAN dynamic reconfiguration *56*
 - subnet *31*
- configure
 - SAN *35*
- contact information *9*

- contacts, Oracle *9*
- control path *16*
- conventions
 - typographical *8*
- customer support *9*

D

- device-mapper
 - function *17*
 - verify installation *37*
- documentation *9*
 - feedback *9*
- download
 - APM software *38*

E

- education programs *9*
- Emulex CNAs
 - settings *24*
- Emulex FCoE
 - supported *24*
- Emulex HBAs
 - install *25*
 - settings *24*
 - supported *24*

F

- FC
 - iSCSI-to-FC routers *34*
 - protocol *23*
 - supported HBAs *23*
 - supported topologies *23*
 - switches *30*
- features
 - APM *10*
 - new in this release *50*
- feedback, documentation *9*
- Fibre Channel
 - Xsigo emulation over InfiniBand *29*
- format
 - LUN disk *44*

H

HBAs

- Brocade *28*
- Emulex *24*
- QLogic *26*

I

install

- APM software *39*
- Brocade CNA drivers *29*
- Brocade CNAs *29*
- Brocade HBA drivers *29*
- Brocade HBAs *29*
- Emulex HBA drivers *25*
- Emulex HBAs *25*
- QLogic CNA drivers *27*
- QLogic CNAs *27*
- QLogic HBA drivers *27*
- QLogic HBAs *27*

iSCSI

how to

- configure Cisco MDS switch *35*
- configure iSCSI initiator *32*
- configure iSCSI-to-FC router *35*
- start iSCSI initiator service *33*

NIC bonding *31*

protocol *23*

supported hardware *30*

supported routers *34*

supported software *30*

supported switches *34*

iSCSI initiator

- start service *33*

iSCSI multipathing

- subnets *31*

iSCSI-to-FC router

- configure *35*

issues

- additional notes *63*
- APM software *51*
- operating system *55*
- Pillar Axiom system *52*

K

KVM hypervisor *13*

L

load balancing

- round-robin *46*
- static *46*

LUN access

- multi-protocol *47*
- precedence *47*

preferred paths *47*

LUN configuration

- SAN hosts *41*

LUN disk

how to

- format *44*
- partition *44*

M

multi-protocol access

- definition *47*
- precedence *47*
- preferred paths *47*

multipath-tools

- function *17*
- pre-installation requirements *22*

N

new features *50*

NIC bonding *31*

non-optimized path *47*

O

online help *9*

operating system

- known issues *55*

optimized path *47*

Oracle Technical Network (OTN) *9*

P

partition

- LUN disk *44*

Pillar Axiom

- configure for LUN access *41*
- interaction with APM *14*

Pillar Axiom system

- known issues *52*
- requirements *20*

pre-configure SAN hosts *35*

pre-installation requirements *22*

product support *9*

Q

QLogic CNAs

- install *27*
- settings *26*
- supported *26*

QLogic HBAs

- install *27*
- settings *26*
- supported *26*

queue depth

- maximum for SAN host HBAs *24*

R

- related documentation *8*
- release notes
 - additional notes *63*
 - APM issues *51*
 - operating system issues *55*
 - Pillar Axiom issues *52*
- remove
 - APM software *49*
- requirements
 - Pillar Axiom system *20*
- round-robin load balancing *46*
- routers
 - iSCSI-to-FC *34*

S

- sales information *9*
- SAN hosts
 - LUN configuration *41*
 - maximum HBA queue depth *24*
 - pre-configure *35*
- SAN protocols *23*
- settings
 - Brocade CNAs *28*
 - Brocade HBAs *28*
 - Emulex CNAs *24*
 - Emulex HBAs *24*
 - QLogic CNAs *26*
 - QLogic HBAs *26*
- static load balancing *46*
- subnets
 - iSCSI multipathing *31*
- Support portal *9*
- supported
 - architectures *19*
 - Brocade CNAs *28*
 - Brocade HBAs *28*
 - Emulex HBAs *24*
 - FC topologies *23*
 - hardware platforms *19*
 - iSCSI hardware *30*
 - iSCSI software *30*
 - Linux kernels *19*
 - QLogic CNAs *26*
 - QLogic HBAs *26*
- switches
 - Cisco MDS *35*
 - FC *30*
 - iSCSI *34*

T

- technical support *9*
- training programs *9*
- typographical conventions *8*

U

- uninstall
 - APM software *49*
- update
 - APM software *39*
 - multipath-tools *39*

V

- virtualization
 - KVM hypervisor *13*
 - Xen hypervisor *13*

X

- Xen hypervisor *13*
- Xsigo I/O Director *29*