

Pillar Axiom Path Manager 3.1



Installation Guide
and Release
Notes

for Red Hat Enterprise Linux 6.1



PILLAR AXIOM

Part Number: 7040340
APM release 3.1
2011 October

Copyright © 2011, 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 © 2011, 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 Pillar Axiom Path Manager

Pillar Axiom Path Manager 3.1 Features.	12
About Pillar Axiom Path Manager and Clustering.	15
About Pillar Axiom Path Manager and Hypervisors.	16
Pillar Axiom Path Manager Architecture.	17
About the Pillar Axiom Path Manager Control Path.	19
About the Pillar Axiom Path Manager Data Path.	20
About Multipathing and Device Mapper Automation.	20
Supported Red Hat Enterprise Linux 6.1 Distributions.	22
Pillar Axiom Path Manager Requirements.	23
Operating Limits.	24

Chapter 2 Install Pillar Axiom Path Manager

Prepare to Install the Pillar Axiom Path Manager.	25
Supported SAN Protocols.	26
Supported Fibre Channel SAN Hardware Components.	26
Supported iSCSI Software and Hardware.	33
Pre-Configure the SAN Host for Pillar Axiom Integration.	38
Verify Installation of the Device Mapper.	40
Download and Install the Pillar Axiom Path Manager Software.	41
Download the Pillar Axiom Path Manager Software.	41
Install (or Update) the Pillar Axiom Path Manager Software.	42
Configure the Pillar Axiom System for LUN Access.	44
Partition and Format the LUN Disk (Optional).	47
About Load Balancing Configurations.	49
About Path Selection.	50
Remove Pillar Axiom Path Manager Software (Optional).	51

Chapter 3 Pillar Axiom Path Manager Release Notes

- New in this Release. 52
- Known APM Issues. 53
- Known Pillar Axiom Issues. 54
- Known Operating System Issues. 61
 - Access to LUNs. 61
 - SAN Dynamic Reconfiguration. 61
 - Multipath Failed Path Errors. 64
 - Non-Optimized Access Events After Slammer Control Unit Failover. 64
 - LUN State Not Updated. 65
 - Too Many Files Open. 65
 - Not All iSCSI Paths Listed. 65
 - Too Many Paths Cause Device-Mapper to Hang. 66
 - Call Trace During Path Failure. 66
 - LUNs Not Shown. 66
 - Manually Load Device Mapper. 67
 - Update Kernel to Show LUNs After Power Off. 67
- Resolved APM Issues. 68
- Additional Notes. 69
 - Mount iSCSI Filesystems. 69
- Index.** 70

List of Figures

Figure 1 APM interaction with a Pillar Axiom server.	18
Figure 2 Example host ports before APM installation.	44
Figure 3 Example host ports after APM installation.	44

List of Tables

Table 1 Typography to mark certain content. 9

Table 2 Contacts at Pillar Data Systems. 10

Table 3 APM 3.1 for RHEL 6.1 features. 12

Table 4 Line and color key for APM interaction diagram. 17

Table 5 Supported hardware platforms. 22

Table 6 APM operating limits. 24

Table 7 Emulex adapter settings. 28

Table 8 QLogic adapter settings. 30

Table 9 Brocade adapter settings. 32

Table 10 Characteristics of SAN connections to host ports. 39

Table 11 Known Pillar Axiom issues. 54

Table 12 Resolved issues. 68

Preface

Audience

This document is intended for individuals who install and maintain Pillar Axiom Path Manager (APM) software.

Expected experience includes:

- Understanding of storage area networks (SANs) and disk storage systems.
- Understanding of Fibre Channel or iSCSI technology.
- Practical knowledge of Oracle Pillar Axiom 600 storage systems.
- Basic Linux administration skills.
- Experience installing software packages on Red Hat Enterprise Linux systems.

Related Documentation

Becoming familiar with other technical documentation (including installation, service, cabling, integration, and administration guides) for Oracle's Pillar Axiom 600 storage system will help you succeed in using this guide.

- *Pillar Axiom Customer Release Notes*: Includes late-breaking important information about the installation and operation of the Pillar Axiom system.
- *Pillar Axiom Administrator's Guide*: Provides detailed information on creating and managing storage resources.
- *Pillar Axiom CLI Reference Guide* (for axiomcli) or *CLI Reference Guide* (for pdscli): Provides detailed information about functions available in the Pillar Axiom command line interfaces (CLIs).

Access Documentation

Technical documentation (including installation, service, cabling, integration, and administration guides) for Oracle’s Pillar Axiom 600 storage system is available from several sources.

- Pillar Axiom Storage Services Manager** Log in to your Pillar Axiom system. Navigate to the **Support** area in the Pillar Axiom Storage Services Manager and select the **Documentation** link.
- Pillar Axiom HTTP access** For Pillar Axiom systems running release 5.0 (and higher) software, point your browser to `http://system-name-IP/documentation.php`, where *system-name-IP* is the name or the public IP address of your system.
- Internet** [Customer support portal](http://support-portal.pillardata.com/csportal/login.seam) (<http://support-portal.pillardata.com/csportal/login.seam>).

Log in and click **Documents** in the left navigation pane.
- Product CD-ROM** Insert the Technical Documentation CD-ROM (came with your Pillar Axiom system) into the CD player and open the DocMenu PDF.

Tip: To search all technical documents on the CD-ROM, click **Search all PDFs** in the top right corner.

Typographical Conventions

Table 1 Typography to mark certain content

Convention	Meaning
<i>italics</i>	Within normal text, words in italics indicate: <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.

Table 1 Typography to mark certain content (continued)

Convention	Meaning
monospace (bold)	<i>Input</i> provided by an administrator on the command line.
>	Indicates a menu item or a navigation path in a graphical user interface (GUI). For example, “Click Storage > Clone LUNs ” means to click the Clone LUNs link on the Storage page in the graphical user interface (GUI).
...	Used within an expression of a navigation path or within a cascading menu structure. The ellipsis indicates that one 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.

Pillar Contacts

Table 2 Contacts at Pillar Data Systems

For help with...	Contact...
Error messages, usage questions, and other support issues	US and Canada: 877-4PILLAR (1-877-474-5527) Europe: +800 PILLAR FS (+800 74 55 27 37) Asia Pacific: +1-408-518-4515 South Africa: +0 800 980 400 Have your system serial number ready. support@pillardata.com Pillar Customer Support (http://support-portal.pillardata.com/csportal/login.seam)
Training (custom or packaged)	Training and Education (http://www.pillardata.com/support-education/training/)
Professional services and inquiries	globalsolutions@pillardata.com

Table 2 Contacts at Pillar Data Systems (continued)

For help with...	Contact...
	Global Solutions (http://www.pillardata.com/support/professional-services/)
Sales and general contact information	Company contacts (http://www.pillardata.com/company/contact)
Documentation improvements and resources	docs@pillardata.com Technical documents (http://www.pillardata.com/techdocs) (Log in with your username and password, and select Documents.)

CHAPTER 1

Introduction to Pillar Axiom Path Manager

Pillar Axiom Path Manager 3.1 Features

Pillar Axiom Path Manager (APM) is defined as:

Optional software installed on a storage area network (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.

The function described in the last bullet enables the Pillar Axiom Storage Services Manager to report information about APM running on the host, such as the number of working paths, and, in some environments, to configure features such as load balancing.

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

Table 3 APM 3.1 for RHEL 6.1 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's Fibre

Table 3 APM 3.1 for RHEL 6.1 features (continued)

Feature	Benefit
	Channel ports, the 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 the Pillar World Wide Customer Support Center. 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 Slammers	Makes connections to Pillar Axiom storage arrays over high-speed FC network infrastructure.
Support for iSCSI connections to both FC and iSCSI Slammers	<p>Makes connections to Pillar Axiom storage arrays over long distances using IP network infrastructure.</p> <p>Note: iSCSI connections to FC Slammers require iSCSI-to-FC routers.</p>
Support for FC clustering	Makes connections to clustered systems operating in an FC-based SAN.
Support for Boot from SAN	Boot from SAN is supported for Fibre Channel SAN.
FCoE CNAs	Fibre Channel over Ethernet (FCoE) Converged Network Adapters (CNAs) are supported on the host.
Support for virtualization	<p>APM supports the KVM hypervisor included in the RHEL 6.1 distribution.</p> <p>The KVM version included in RHEL 6.1 is supported on x86-64 and not on x86-32.</p>

Table 3 APM 3.1 for RHEL 6.1 features (continued)

Feature	Benefit
Support for load balancing	Both static and round-robin load balancing options are supported.

Related concepts

- [About Pillar Axiom Path Manager and Clustering](#)
- [About Pillar Axiom Path Manager and Hypervisors](#)

About Pillar Axiom Path Manager and Clustering

The Pillar Axiom Path Manager (APM) for Red Hat Enterprise Linux 6.1 (RHEL 6.1) can be used in a clustered environment that uses Fibre Channel connections.

The APM software must be installed and working before you set up clustering.

If the cluster function fails on the nodes, but the nodes are still running at the operating system level, data consistency can be affected. The leftover write operations from failed database instances reach the Pillar Axiom storage system after the recovery process starts. Because these write operations are no longer in the proper serial order, they can damage the consistency of the stored data. To avoid damaging the data consistency, when a cluster node fails, fence the failed node off from all the shared disk devices or disk groups.

This I/O fencing process is sometimes called *I/O fencing-exclusion*, *disk fencing*, or *failure fencing*. The purpose of I/O fencing is to:

- Prevent updates by failed instances.
- Detect failures.
- Prevent the *split-brain* effect in the cluster.

Note: The split-brain effect occurs when failure of a single cluster causes reconfiguration of the cluster into multiple partitions, each of which assumes ownership of the same exclusive resources, usually resulting in data corruption.

The Cluster Logical Volume Manager (CLVM), in association with the shared storage unit, and the Red Hat Global File System, are required to prevent the failed nodes from accessing shared devices.

For more information on how to set up clustering in your environment, refer to the RHEL 6 [documentation](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Cluster_Administration/index.html) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Cluster_Administration/index.html).

Note: The RHEL Cluster Suite installs and enables the `scsi_reserve` service by default. This service should be disabled if the `scsi_fence` option is not used as a fencing mechanism. To disable the `scsi_reserve` service, run the following command:

```
# chkconfig scsi_reserve off
```

About Pillar Axiom Path Manager and Hypervisors

Pillar Axiom Path Manager (APM) supports the KVM hypervisor open source virtual machine monitor included with the Red Hat Enterprise Linux 6.1 (RHEL 6.1) distribution on x86-64 platforms but not on x86-32.

You can install APM on a RHEL 6.1 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.

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.

For information on using disks with the KVM hypervisor, refer to the [Red Hat Enterprise Linux 6 Virtualization Guide](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html-single/Virtualization/index.html) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html-single/Virtualization/index.html).

Pillar Axiom Path Manager Architecture

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

[Figure 1: APM interaction with a Pillar Axiom server](#) illustrates how the APM software installed on a storage area network (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-supplied hardware and software
	Non-Pillar hardware and software
	SAN host kernel space
	SAN host user space

Figure 1 APM interaction with a Pillar Axiom server



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 paths (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 server
9 Network card	18 Brick storage enclosure pool

Related concepts

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

About the Pillar Axiom Path Manager 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 running in the background to control multipathing and communication. The APM daemon uses the control path to:

- Get Fibre Channel (FC) and Internet Small Computer System Interface (iSCSI) port information from the host bus adapter (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 host's iSCSI initiator.

The graphical user interface (GUI) and command line interface (CLI) list the port World Wide Names (WWNs) of the FC ports in the host and the Internet Protocol (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. Within APM, the connections appear to originate from the FC ports that are assigned on the switch to the host's iSCSI initiator. The WWNs of these ports are displayed as Fibre Channel 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 its Pilot to the APM host over the data path.

About the Pillar Axiom Path Manager 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.

See [Figure 1](#) for an illustration of how data flows from the host to the Pillar Axiom system.

The Linux multipath framework:

- 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 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.
- Manages data path errors.

An optimized path provides the best performance and is the preferred path for data transfer.

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 plug-ins for the `multipath-tools` daemon that monitor the health of paths to Pillar Axiom LUNs, and determine the appropriate priority of each path. APM also provides a daemon that monitors the state of the `device-mapper` and communicates with the Pillar Axiom software. APM presents the multipathed LUNs as virtual block devices in the Linux 2.6 `device-mapper` framework.

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 usually required by `device-mapper`. 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 Red Hat Enterprise Linux 6.1 Distributions

Pillar Axiom Path Manager (APM) is supported on Red Hat Enterprise Linux 6.1 (RHEL 6.1) platforms.

APM supports RHEL 6.1 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	Output from <code>uname -i</code>
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:

```
2.6.32-71
```

Pillar Axiom Path Manager Requirements

Pillar Axiom systems presenting LUNs to Red Hat Enterprise Linux 6.1 (RHEL 6.1) hosts using Pillar Axiom Path Manager (APM) 3.1 must be running release 3.5.6 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, the `device-mapper`, Linux, 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
Handle data paths	32 to each LUN
Handle FC HBA ports	32 for each SAN host

The Linux `device-mapper` has a limitation of 1024 paths to all devices, including Pillar Axiom LUNs. If you use the maximum of 256 LUNs, you will have a maximum of four paths to each LUN.

Important! The Linux `device-mapper` does not gracefully handle more than 1024 visible paths and may fail in a variety of ways if that limit is exceeded.

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 Pillar Axiom Path Manager

Prepare to Install the Pillar Axiom Path Manager

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

- 1 Read the Pillar Axiom Path Manager Release Notes in Chapter 3 of this document.
- 2 Ensure that the Pillar Axiom system is running release 3.5.6 or higher of the Pillar Axiom software.
- 3 If you are using Fibre Channel (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.

The following types of adapters are supported:

- Emulex
 - QLogic
 - Brocade
- 5 If you are using Internet Small Computer System Interface (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](#)

Related references

- [Supported Fibre Channel SAN Hardware Components](#)
- [Supported iSCSI-to-Fibre Channel Routers](#)

Related tasks

- [Install Emulex Fibre Channel Adapters and Drivers](#)
- [Install QLogic Fibre Channel Adapters and Drivers](#)
- [Install Brocade Fibre Channel 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) 3.1 supports Fibre Channel (FC) and software-based Internet Small Computer System Interface (iSCSI) connections to the Pillar Axiom system.

Supported Fibre Channel SAN Hardware Components

Pillar Axiom Path Manager (APM) supports a variety of Fibre Channel (FC) host bus adapters (HBAs), converged network adapters (CNAs), and storage area network (SAN) routers and switches.

Approved HBAs, CNAs, and switches that are compatible with the Pillar Axiom system are listed in the following related references.

Related references

- [Supported Emulex Fibre Channel Adapters and Drivers](#)
- [Supported QLogic Fibre Channel Adapters and Drivers](#)
- [Supported Brocade Fibre Channel Adapters and Drivers](#)
- [Supported Fibre Channel and Fibre Channel over Ethernet Switches](#)

Supported Fibre Channel 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 may cause errors.

Supported Emulex Fibre Channel Adapters and Drivers

Pillar Axiom Path Manager (APM) supports the following Emulex Fibre Channel (FC) host bus adapters (HBAs) and HBA drivers:

- 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 FCoE CNAs can also be used concurrently as Ethernet network interface controllers (NICs).

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

Related references

- [Recommended Settings for Emulex Fibre Channel Adapters](#)

Related tasks

- [Install Emulex Fibre Channel Adapters and Drivers](#)

Recommended Settings for Emulex Fibre Channel Adapters

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

Table 7 Emulex adapter settings

Adapter setting	Recommended value
<code>lpfc_lun_queue_depth</code>	64
<code>lpfc_devloss_tmo</code>	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 Fibre Channel Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your Emulex host bus adapters (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 a 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 force after the next reboot.

Supported QLogic Fibre Channel Adapters and Drivers

The Pillar Axiom Path Manager (APM) software supports the following QLogic Fibre Channel (FC) host bus adapters (HBAs) on x86-32 and x86-64 platforms:

- 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 FCoE CNAs can also be used concurrently as Ethernet network interface controllers (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 Fibre Channel Adapters](#)

Related tasks

- [Install QLogic Fibre Channel Adapters and Drivers](#)

Recommended Settings for QLogic Fibre Channel Adapters

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

Table 8 QLogic adapter settings

Adapter setting	Recommended value
<code>ql2xloginretrycount</code>	30
<code>ql2xlogintimeout</code>	60
<code>qlport_down_retry</code>	45
<code>ql2xmaxqdepth</code>	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. See the QLogic website for details.

Install QLogic Fibre Channel Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your QLogic host bus adapters (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://support.qlogic.com/support) (currently <http://support.qlogic.com/support>).
- 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 a 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 picked up.

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 Fibre Channel Adapters and Drivers

Pillar Axiom Path Manager (APM) supports the following Brocade Fibre Channel (FC) host bus adapters (HBAs):

- 415
- 425
- 815
- 825

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

- 1010
- 1020

Note: These FCoE CNAs can also be used concurrently as Ethernet network interface controllers (NICs).

These HBAs and CNAs require the 3.0.0.0 (or later) driver, available from the [Brocade website](http://www.brocade.com/sites/dotcom/services-support/drivers-downloads/CNA/Linux.page) (<http://www.brocade.com/sites/dotcom/services-support/drivers-downloads/CNA/Linux.page>).

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

Related references

- [Recommended Settings for Brocade Fibre Channel Adapters](#)

Related tasks

- [Install Brocade Fibre Channel Adapters and Drivers](#)

Recommended Settings for Brocade Fibre Channel Adapters

We recommend the following settings for Brocade host bus adapters (HBAs) and converged network adapters (CNAs):

Table 9 Brocade adapter settings

Adapter setting	Recommended value
pathtov	60

Install Brocade Fibre Channel Adapters and Drivers

To ensure compatibility between Pillar Axiom Path Manager (APM) and your Brocade host bus adapters (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 download page](http://www.brocade.com/services-support/drivers-downloads/CNA/Linux.page#Docs) (currently <http://www.brocade.com/services-support/drivers-downloads/CNA/Linux.page#Docs>).
- 2 Download version 3.0.0.0 or later of the driver and utilities package from the [Brocade support download page](#).
- 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 Fibre Channel (FC) and FC over Ethernet (FCoE) port.
The path time out value should be set to 60. See the Brocade documentation for details of how to set the value.

Supported Fibre Channel and Fibre Channel over Ethernet Switches

For a list of supported Fibre Channel (FC) and FC over Ethernet (FCoE) switches, choose one of:

- Call the Pillar World Wide Customer Support Center at the number listed in the Pillar Contacts related reference.
- Refer to the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Technical Documents](http://www.pillardata.com/techdocs) website (<http://www.pillardata.com/techdocs>).

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

Related references

- [Pillar Contacts](#)

Supported iSCSI Software and Hardware

Pillar Axiom Path Manager (APM) for Red Hat Enterprise Linux 6.1 (RHEL 6.1) supports the following Internet Small Computer System Interface (iSCSI) software and hardware:

- The iSCSI software initiator included in the RHEL 6.1 distribution.
- The iSCSI-to-Fibre Channel routing features of the Cisco MDS 9000 family of routers and switches.

Note: APM for RHEL 6.1 does not support iSCSI host bus adapters (HBAs).

About Network Configuration for iSCSI Multipathing

You can reduce the impact of failures in your Internet Small Computer System Interface (iSCSI) network by ensuring that iSCSI connections are made through several network interface cards (NICs) on the host.

We recommend using the normal IP configuration facilities to route connections to different Slammer storage controller ports 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 network interface cards (NICs) and Slammer ports can be configured into different subnets. Whenever the Internet Small Computer System Interface (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 Internet Protocol (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 Transmission Control Protocol/Internet Protocol (TCP/IP), Internet Small Computer System Interface (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) for Red Hat Enterprise Linux 6.1 (RHEL 6.1) requires the version of the iSCSI software initiator included in the RHEL 6.1 distribution.

For complete installation and configuration instructions, refer to: `/usr/share/doc/iscsi-initiator-utils-*/README`. The following procedure is 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 may 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.

```
# 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 Discover the other ports, depending on the desired design (optional).

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

```
# 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 Internet Small Computer System Interface (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-Fibre Channel Routers

iSCSI-to-Fibre Channel (FC) routing features enable a host to use Internet Small Computer System Interface (iSCSI) to access LUNs on Pillar Axiom FC 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-Fibre Channel Router](#)

Supported iSCSI Switches

For a list of supported iSCSI switches, you can:

- Call the Pillar World Wide Customer Support Center.
- Refer to the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Customer support portal](http://support-portal.pillarata.com/csportal/login.seam) (<http://support-portal.pillarata.com/csportal/login.seam>).

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

Related references

- [Pillar Contacts](#)

Configure the iSCSI-to-Fibre Channel Router

The iSCSI-to-Fibre Channel 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 Pillar Axiom Slammer storage controller ports as Internet Small Computer System Interface (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 Fibre Channel 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, you need to prepare your storage area network (SAN) host connections with the Pillar Axiom system.

Prerequisites:

Verify that your system has:

- At least one SAN Slammer storage controller that has Fibre Channel (FC) or Internet Small Computer System Interface (iSCSI) ports.
- Supported host bus adapter (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.

Tip: To check network connectivity, issue a simple `pdcli` or `axiomcli` request from the host to the Pillar Axiom system. Both `pdcli` and `axiomcli` use the same port and protocols as those used by APM. Refer to the *Pillar Axiom CLI Reference Guide* or *Pillar Axiom CLI Reference Guide* for details.

- 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 server.
- 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

See [Figure 2: Example host ports before APM installation](#) for an illustration.

Related concepts

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

Related tasks

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

Verify Installation of the Device Mapper

The Linux 2.6 `device-mapper`, among other uses, maps a logical block device on the SAN host into 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.

Download and Install the Pillar Axiom Path Manager Software

The Pillar Axiom Path Manager (APM) installation for Red Hat Enterprise Linux 6.1 (RHEL 6.1) requires that you download the APM software package from the customer support portal, and install the software on your system.

- 1 Download the APM software.
- 2 Install the APM software.
- 3 Configure the Pillar Axiom software for LUN access.

To remove the APM software from your SAN host, see the related task.

Related tasks

- [Download the Pillar Axiom Path Manager Software](#)
- [Install \(or Update\) the Pillar Axiom Path Manager Software](#)
- [Remove Pillar Axiom Path Manager Software \(Optional\)](#)

Download the Pillar Axiom Path Manager Software

Download the Pillar Axiom Path Manager (APM) software from the customer support portal.

The APM software consists of a single package that you will need to download.

- 1 Log in to the [customer support portal](http://support-portal.pillardata.com/csportal/login.seam) (<http://support-portal.pillardata.com/csportal/login.seam>).
- 2 In the left-hand navigation pane, click **Software Downloads > Pillar Axiom Path Manager**.
- 3 In the right-hand content pane, navigate to the name of the installation package for your hardware platform.
- 4 Click the name of the package for your hardware platform to download (x86 or AMD64/Intel EM64T).
- 5 In the **Software Download Details** pane below, click the green arrow, and follow the prompts.
- 6 To download the package to your SAN host, choose the **Save** option.

Install (or Update) the Pillar Axiom Path Manager Software

After you download the Pillar Axiom Path Manager (APM) software you can install it on your host system.

Before you install the APM software, verify that your system meets all hardware and software prerequisites.

Two steps in the installation process enable you to keep your original multipath configuration file (`/etc/multipath.conf`) so that you can continue to use 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 `/etc/multipath.conf`.

- 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 using the following command:

```
# rpm -e axiompmp
```

- 3 Install the APM package using 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.

After you install the software, configure your Pillar Axiom system for LUN access. You must complete that task for the software to function correctly.

Related references

- [*Supported Fibre Channel SAN Hardware Components*](#)

Related tasks

- [*Pre-Configure the SAN Host for Pillar Axiom Integration*](#)
- [*Configure the Pillar Axiom System for LUN Access*](#)

Configure the Pillar Axiom System for LUN Access

The LUN configuration procedure provides the storage area network (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 software.
- 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 server for the SAN hosts.

Note: In a clustered environment, LUNs shared in a cluster should be mapped to all host nodes in the cluster.

- 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 Run the following commands to list the multipath devices (optional):

```
# /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.

- 10 If you plan to do so, follow the instructions to partition and format the LUN disk on the SAN host.

Related references

- [*SAN Dynamic Reconfiguration*](#)

Related tasks

- [*Pre-Configure the SAN Host for Pillar Axiom Integration*](#)
- [*Install \(or Update\) the Pillar Axiom Path Manager 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` 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/axiompmd 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 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*.

About Path Selection

Pillar Axiom Path Manager (APM) supports access to LUNs using Internet Small Computer System Interface (iSCSI) and Fibre Channel (FC) protocol at the same time, as well as individually.

Paths to a LUN may have different performance characteristics. Paths to the Slammer control unit (CU) where the LUN resides are considered *optimized paths*; other paths to the LUN are considered *non-optimized paths*. When both FC and iSCSI access to a LUN are available, FC access generally performs better.

APM divides the paths to each LUN into four groups with different performance characteristics, in this order of precedence:

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

At any given time, the `device-mapper` framework only uses paths from the most preferred group that has paths available.

Note: When an active path fails, and I/O traffic is transferred to a different path, I/O performance will be reduced for a short time while the operating system recognizes the failure and makes the path transfer. If the failing path was optimized and the new path is non-optimized, I/O performance may remain lower than before since it is now using a lower-performance path. Within a few minutes of traffic being transferred to a non-optimized path, the Pillar Axiom system will attempt to reconfigure the LUNs to use an optimized path, and upon success, I/O performance will improve.

Remove Pillar Axiom Path Manager Software (Optional)

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

Before you uninstall APM, if you do not want to access Pillar Axiom LUNs, we recommend that you disconnect the storage area network (SAN) host from the Pillar Axiom system.

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

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

After the APM software 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 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), and you may need to reconfigure the host to access the LUNs through new device names.

CHAPTER 3

Pillar Axiom Path Manager Release Notes

New in this Release

- Pillar Axiom Path Manager (APM) 3.1 supports Red Hat Enterprise Linux 6.1 (RHEL 6.1) 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>When a Fibre Channel HBA is removed from a host running APM, it remains associated with that host.</p> <p>If the HBA is moved to a host that is either not running APM or on which APM is shown as not communicating with the Pillar Axiom Pilot, any LUNs mapped to the host will continue to be accessible through the HBA ports. The GUI and CLI will continue to report the HBA as being present in the original host.</p> <p>If the HBA is moved to a host where APM is running and communicating with the Pilot, its association and mappings for the old host will be removed, and the mappings for the new host will be applied.</p>	<p>This issue is fixed in release 4.0 of the Pillar Axiom software.</p>
<p>If you use the GUI or CLI to change the default configured Slammer control unit (CU) of a LUN to the other CU on the Slammer, the Slammer port mask for the LUN will be reversed. For example, if CU0 Port 0 is enabled and Port 1 is excluded, and the LUN is moved to CU1, then CU1 Port 0 will become excluded and CU1 Port 1 will become enabled.</p>	<p>This issue is fixed in release 4.0.0 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.0.0, after you change the default configured Slammer CU for a LUN, be sure to update the LUN port mask to the required value.</p>
<p>If an APM host uses iSCSI to connect to a Pillar Axiom system, and it uses an iSCSI initiator name that is the same as its host name, then the entry for that host in the Pillar Axiom Storage Services Manager will be continually deleted and recreated, causing the host</p>	<p>This issue is fixed in release 4.1 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.1, ensure that the iSCSI initiator names configured on hosts that use iSCSI to connect to a</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
<p>entry to disappear and reappear intermittently.</p>	<p>Pillar Axiom system are different from all host names used by APM hosts, including virtual machine hosts, connected to that Pillar Axiom system. The iSCSI standards require that iSCSI names follow particular formats, as specified in RFC 3720: (http://tools.ietf.org/html/rfc3720#section-3.2.6).</p> <p>If hosts are configured to use iSCSI initiator names that conform to these requirements, it is extremely unlikely that they will be the same as any host name.</p>
<p>If all paths to a LUN's configured Slammer control unit (CU) fail, APM will re-route all traffic through the non-optimized paths to the LUN's alternate CU. In response, the Pillar Axiom system will initially log events indicating non-optimized access, then when this traffic continues it will temporarily move the LUN to the alternate CU. This process leaves the host using optimized paths to the LUN, but the LUN is resident on a CU other than its configured home.</p> <p>Normally, the system will attempt to move the LUN back to its configured CU from time to time, and if the paths to the other CU have recovered the traffic will transfer back and the system returns to its normal configured state. However, if the Pilot software is restarted while a LUN is in this temporary state, as might happen during a software update that includes the option to update the Pilot software, two problems occur:</p> <ol style="list-style-type: none"> 1 The graphical user interface (GUI) and command line interface (CLI) wrongly report that the 	<p>This issue is fixed in release 4.0 of the Pillar Axiom software.</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
<p>LUN's current CU is its configured CU.</p> <p>2 Non-optimized access events are no longer logged for the LUN, and the system does not attempt to move the LUN back to its configured CU.</p> <p>If subsequent path failures and recoveries cause traffic to be sent to the CU on which the LUN is not resident, the system will not move the LUN to the CU receiving the traffic. This means that all traffic to the LUN would have non-optimized access, which decreases performance, and this non-optimized access would not be logged.</p>	
<p>When a LUN is created on a Pillar Axiom system, its load balancing attribute is set to round-robin by default. If the LUN is then mapped to a host running APM, the load balancing attribute setting can change to static when APM on the host first communicates with the Pillar Axiom system after detecting the LUN. Instead, this attribute should be set to round-robin when the LUN is first created, and should change only if an administrator changes it using the GUI or CLI.</p>	<p>This issue is fixed in Pillar Axiom software release 4.0.</p> <p>If the Pillar Axiom system is running a release earlier than 4.0, check that the load balancing attribute for the LUN is still set to the desired value after APM on the host has detected the LUN and its LUN name on Host has been reported in the GUI. If the setting has changed, change it back to the desired value, which can then be correctly saved.</p>
<p>If an iSCSI initiator is added to a SAN host that has authentication enabled, authentication will not be enabled for that initiator. Also, if an iSCSI initiator with authentication enabled on a previous SAN host is moved to another host with or without authentication enabled, the initiator will retain its original configuration.</p>	<p>This issue is fixed in Pillar Axiom software release 4.0.</p> <p>If the Pillar Axiom system is running a release earlier than 4.0, disable then re-enable authentication for the host after iSCSI initiators are added to or moved between SAN hosts.</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
The GUI and CLI may show incorrect link speeds for SAN hosts with 8 Gb/s HBAs.	This issue is fixed in release 4.1 of the Pillar Axiom software.
After recovery from a Slammer control unit (CU) failure, the Pillar Axiom system may become incapable of automatically moving LUNs between the CUs on that Slammer. When the system attempts to move the LUNs automatically in response to non-optimized access from a host, the attempts fail, and non-optimized access persists.	<p>This issue is fixed in release 4.1 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.1, use the GUI or CLI to re-assign the LUNs to the CUs through which access is currently taking place. Alternatively, restart the Pillar Axiom system to restore optimized access.</p>
When an iSCSI initiator name is changed or removed on an APM host, the GUI and CLI may continue to associate the old name with the host.	<p>This issue is fixed in release 4.2 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.2:</p> <ol style="list-style-type: none"> 1 Stop the APM on the host. 2 When the Pillar Axiom system reports the host as Not Connected, delete the host. <p style="padding-left: 40px;">This will preserve LUN mappings to the initiators.</p> <ol style="list-style-type: none"> 3 Restart the APM daemon.
<p>When a Pilot restart occurs on a Pillar Axiom system running release 4.0 (4.0.4 or later) or release 4.1 (4.1.0 or later) of the Pillar Axiom software, all LUNs on the system move from their current Slammer control unit (CU) to the other Slammer CU. As a result, the optimization of all paths to the LUNs changes.</p> <p>Note: A non-disruptive upgrade to release 4.0 (4.0.4 or later) or release 4.1 (4.1.0 or later) of the Pillar Axiom software will cause a Pilot restart, which will trigger this problem.</p>	<p>This issue is fixed in release 4.1.4 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.1.4, and if all relevant SAN hosts have paths to both CUs on the Slammers, and those paths are managed by an ALUA-aware path management system such as Pillar Axiom Path Manager, it should not be necessary to take any action. The LUNs will remain balanced across the CUs, and the path management software will ensure that only optimized paths to the LUNs are used. Be aware that traffic may be moved to alternate</p>

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
	paths when a Pilot restart occurs. Other hosts may need their path configuration to be changed to ensure that they access each LUN through its new current home CU. Alternatively, all LUNs can be moved back to their default configured CU by restarting the Pillar Axiom system.
The load balancing configuration displayed in the GUI or CLI can be different from the load balancing setting that the Pillar Axiom system sends to APM on the SAN host. This happens because, occasionally in the course of updating its configuration records to describe the new host, the Pillar Axiom system creates duplicate internal records.	This issue is fixed in Release 4.2.3 and later of the Pillar Axiom software.
A change in the number of reported iSCSI initiator IQNs can cause the loss of all SAN zonings. If an APM host connects to the Pillar Axiom system using more than one iSCSI Initiator name, and is subsequently reconfigured to have fewer iSCSI Initiator names, some or all LUN mappings over iSCSI to that host may be lost. The Pillar Axiom Storage Services Manager will still show iSCSI mappings to the host, but these mappings will not be activated and the host will be unable to access the LUNs over some or all iSCSI connections.	This issue is fixed in release 4.3 of the Pillar Axiom software. If the Pillar Axiom system is running a release earlier than 4.3, use the Axiom GUI or CLI to delete the mappings to the host for the affected LUNs, and then recreate the mappings.
In the GUI, LUN host connections in a disconnected state may display Protocol Disabled instead of Not Connected status in the View Connections screen.	This issue is fixed in release 4.3 of the Pillar Axiom software.
When a LUN connection protocol has been disabled, the connection status for	This issue is fixed in release 4.0 of the Pillar Axiom software.

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
the LUN that appears in the Pillar Axiom GUI will be invalid.	If the Pillar Axiom system is using a release earlier than 4.0, you can safely ignore the protocol-specific connection status of any LUN for which that connection protocol has been disabled.
If an iSCSI initiator is connected to port on a Slammer, and that Slammer Control Unit is powered off, the LUN and Host GUI pages continue to show the connection status for the iSCSI initiator as <code>Connected</code> .	This issue is fixed in release 5.0 of the Pillar Axiom software.
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. When this happens, the control path connection between APM and the Axiom will continually move between <code>Communicating</code> and <code>Not Communicating</code> 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.	This issue is fixed in release 5.0 of the Pillar Axiom software.
When a Fibre Channel (FC) HBA is replaced or removed from an APM host, and the Pillar Axiom system is running release 4.0 or later, the LUN mappings are removed from all FC ports in the host. The mappings still appear in the GUI to be mapped to the host but they are not actually mapped to the host's FC ports.	<p>This issue is fixed in Release 4.3.9 of the Pillar Axiom software.</p> <p>If the Pillar Axiom system is running a release earlier than 4.3.9:</p> <ol style="list-style-type: none"> 1 Before removing the old HBAs, stop the APM daemon and delete the host entry in the Pillar Axiom system. 2 Create a temporary host entry in the Pillar Axiom system and associate the old HBAs with it. 3 Make the changes to the HBAs in the host.

Table 11 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
	<p>4 Add all the HBAs now in the host to the temporary host entry created in step 2.</p> <p>5 Start the APM daemon on the host. Choose one of:</p> <ul style="list-style-type: none"> • Reboot the host. • Restart the daemon if the host was not shut down. <p>Once the daemon connects to the Pillar Axiom system, it will recreate its host entry and transfer the mappings from the temporary entry created in step 2.</p> <p>6 Delete the temporary host entry created in step 2.</p>
<p>The Pillar Axiom Storage Services Manager (GUI) sometimes continues to display host paths that are no longer valid after APM stops. 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>
<p>If you notice a decrease in I/O performance between a LUN and a Slammer control unit (CU), it might be the result of a problem recovering from temporary use of non-optimized paths to the alternate CU on that Slammer.</p>	<p>This issue will be fixed in a future release of the Pillar Axiom software.</p> <p>To restore the optimized paths, verify your SAN connections. Then follow the instructions in the <i>Pillar Axiom Administrator's Guide</i> to rehome the LUN to the alternate CU on that Slammer. After that, you will need to rehome the LUN again to the original CU on that Slammer. Contact the Pillar World Wide Customer Support Center for assistance.</p>

Known Operating System Issues

The following operating system issues may have an impact on running APM on systems.

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 \(Optional\)](#)

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/ch23.html) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch23.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.

If, after creating a new LUN on the Pillar Axiom system and assigning it to the host, running the rescan script does not bring up the LUNs, this may have happened because Linux incorrectly believes that the LUN number is 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.

- The Emulex HBA driver for Linux enables you to dynamically add or delete LUNs and targets without unloading or reloading the `lpfc` module and without resetting the adapter. Use the Emulex `lun_scan` script in `/usr/sbin`. Refer to the [Emulex Driver for Linux User Manual](http://www-dl.emulex.com/support/linux/82022/manual.pdf) (<http://www-dl.emulex.com/support/linux/82022/manual.pdf>) for details.
- For QLogic HBAs, use the Dynamic Target and LUN Discovery script `ql-dynamic-tgt-LUN-disc.sh`, available from the [QLogic Downloads page](http://driverdownloads.QLogic.com) (<http://driverdownloads.QLogic.com>).

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.

After LUN deletion, you may see a message similar to the following while the Pillar Axiom system is restarting the daemon:

```
error calling out /sbin/scsi_id -g -u -s /block/sd*
```

This message indicates that entries for the deleted LUNs still exist in the `device-mapper` device table. Rebooting the host will flush these 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://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch23.html) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Storage_Administration_Guide/ch23.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.

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.

Non-Optimized Access Events After Slammer Control Unit Failover

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, or after all paths to a CU fail and subsequently recover.

After a CU failover and failback, 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.

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.

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.

Not All iSCSI Paths Listed

On hosts with a large number of devices, all iSCSI paths may not be listed.

This is the result of a problem in the Linux iSCSI and `device-mapper` code.

To display the missing paths, first rescan all iSCSI devices using the following command:

```
iscsiadm -m node -R
```

Then restart the `multipathd` service using the following command:

```
service multipathd restart
```

Too Many Paths Cause Device-Mapper to Hang

Due to a Linux defect, the `device-mapper` daemon occasionally hangs when there are more than 512 LUNs or device paths.

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
```

The operating system vendor will provide an update to fix this defect.

Call Trace During Path Failure

Due to a Linux defect, call trace messages can appear during path failure.

These messages are only informational and do not disrupt I/O. It is not a Pillar Axiom Path Manager (APM) problem.

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.

The operating system vendor will provide a fix in future operating system updates.

Manually Load Device Mapper

It may be necessary to load the `device-mapper` driver manually after Pillar Axiom Path Manager (APM) installation.

If the error message `DM multipath kernel driver not loaded` is displayed during APM installation, it may be necessary to load the `device-mapper` driver manually using the following command:

```
modprobe dm_multipath
```

This is normal Linux behavior because the Linux operating system does not automatically load the `device-mapper` driver upon restarting the `multipathd` service.

Update Kernel to Show LUNs After Power Off

Due to a Linux defect, after a Slammer control unit (CU) is powered off, some LUNs may not appear in the output of the `/sbin/multipathd show maps topology` command.

Red Hat has provided a fix for this issue.

- 1 Use the `yum` tool to update your Red Hat kernel to version 2.6.32-131.12.1 or later.
- 2 Before you run the `/sbin/multipathd show maps topology` command again, run the following command to show all LUNs and paths:

```
/sbin/multipath -v3
```

Resolved APM Issues

The issues listed in the following table have been resolved in the current release of Pillar Axiom Path Manager (APM).

Table 12 Resolved issues

Issue
The Pillar Axiom Path Manager (APM) daemon reports an unknown port speed entry for Brocade CNAs due to a bug in the Brocade HBA API.
The name and path information for LUNs using friendly names is not displayed in the Pillar Axiom GUI .

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 *69*

APM

control path *19*

data path *20*

definition *12*

features *12*

how to

download software *41*

install software *42*

remove software *51*

update software *42*

in a clustered environment *15*

interaction with Pillar Axiom *17*

known issues *53*

operating limits *24*

pre-installation requirements *25*

audience *8*

B

Brocade CNAs

install *32*

settings *32*

supported *31*

Brocade HBAs

install *32*

settings *32*

supported *31*

C

Cisco MDS switch

configure *37*

clustered environment *15*

disk fencing *15*

failure fencing *15*

I/O fencing *15*

split-brain *15*

CNAs

Brocade *31*

QLogic *29*

configuration

iSCSI multipathing *33*

iSCSI-to-FC router *37*

Pillar Axiom LUNs *44*

SAN *38*

SAN dynamic reconfiguration *61*

subnet *34*

contact information *10*

conventions

typographical *9*

D

device-mapper

function *20*

verify installation *40*

disk fencing *15*

documentation

accessing *9*

suggestions *11*

download

APM software *41*

multipath-tools *41*

E

education programs *10*

Emulex CNAs

settings *28*

Emulex FCoE

supported *27*

Emulex HBAs

install *28*

settings *28*

supported *27*

experience

required *8*

F

failure fencing *15*

features

APM *12*

new in this release *52*

Fibre Channel

iSCSI-to-FC routers *37*

protocol *26*

- supported HBAs *26*
- supported topologies *27*
- switches *32*

- format
 - LUN disk *47*

H

HBAs

- Brocade *31*
- Emulex *27*
- QLogic *29*

help

- online *10*

I

I/O fencing *15*

install

- APM software *42*
- Brocade CNA drivers *32*
- Brocade CNAs *32*
- Brocade HBA drivers *32*
- Brocade HBAs *32*
- Emulex HBA drivers *28*
- Emulex HBAs *28*
- multipath-tools *42*
- QLogic CNA drivers *30*
- QLogic CNAs *30*
- QLogic HBA drivers *30*
- QLogic HBAs *30*

iSCSI

how to

- configure Cisco MDS switch *37*
- configure iSCSI initiator *35*
- configure iSCSI-to-FC router *37*
- start iSCSI initiator service *36*

NIC bonding *34*

- protocol *26*
- supported hardware *33*
- supported routers *37*
- supported software *33*
- supported switches *37*

iSCSI initiator

- start service *36*

iSCSI multipathing

- subnets *34*

iSCSI-to-FC router

- configure *37*

issues

- additional notes *69*
- APM *53*
- operating system *61*
- Pillar Axiom *54*
- resolved *68*

K

- KVM hypervisor *16*

L

load balancing

- round-robin *49*
- static *49*

LUN access

- multi-protocol *50*
- precedence *50*
- preferred paths *50*

LUN configuration

- SAN hosts *44*

LUN disk

- how to
 - format *47*
 - partition *47*

M

multi-protocol access

- definition *50*
- precedence *50*
- preferred paths *50*

multipath-tools

- function *20*
- how to
 - download software *41*
 - install software *42*
 - remove software *51*
 - update software *42*
- pre-installation requirements *25*

N

new features *52*

NIC bonding *34*

non-optimized path *50*

O

online documents *9*

online help *10*

operating system

- known issues *61*

optimized path *50*

P

partition

- LUN disk *47*

Pillar Axiom

- configure for LUN access *44*
- interaction with APM *17*
- known issues *54*
- software requirements *23*

Pillar Data Systems

- Support portal *10*

pre-configure SAN hosts *38*

pre-installation requirements *25*

product support *10*
professional services *10*

Q

QLogic CNAs
 install *30*
 settings *30*
 supported *29*

QLogic HBAs
 install *30*
 settings *30*
 supported *29*

queue depth
 maximum for SAN host HBAs *27*

R

release notes
 additional notes *69*
 APM issues *53*
 operating system issues *61*
 Pillar Axiom issues *54*
 resolved issues *68*

remove
 APM software *51*
 multipath-tools *51*

requirements
 Pillar Axiom software *23*

requisite reading *8*
resolved issues *68*
round-robin load balancing *49*

routers
 iSCSI-to-FC *37*

S

sales information *11*
SAN hosts
 LUN configuration *44*
 maximum HBA queue depth *27*
 pre-configure *38*

SAN protocols *26*

settings
 Brocade CNAs *32*
 Brocade HBAs *32*
 cluster timeout *15*
 Emulex CNAs *28*
 Emulex HBAs *28*
 QLogic CNAs *30*
 QLogic HBAs *30*

solutions (professional services) *10*

split-brain *15*
static load balancing *49*

subnets
 iSCSI multipathing *34*

Support portal *10*

supported
 architectures *22*
 Brocade CNAs *31*
 Brocade HBAs *31*
 Emulex HBAs *27*
 Fibre Channel topologies *27*
 hardware platforms *22*
 iSCSI hardware *33*
 iSCSI software *33*
 Linux kernels *22*
 QLogic CNAs *29*
 QLogic HBAs *29*
switches
 Cisco MDS *37*
 Fibre Channel *32*
 iSCSI *37*

T

technical documents
 accessing *9*
technical support *10*
training programs *10*
typographical conventions *9*

U

uninstall
 APM software *51*
 multipath-tools *51*
update
 APM software *42*
 multipath-tools *42*

V

virtualization
 KVM hypervisor *16*
 Xen hypervisor *16*

X

Xen hypervisor *16*