

Pillar Axiom Path Manager 3.2



Installation Guide
and Release
Notes

for Oracle VM Server 2.2 for x86



PILLAR AXIOM

Part Number: E29499-01

APM release 3.2

2012 May

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Table of Contents

Preface

Chapter 1 Introduction to APM

APM Features.	10
About APM and Xen Hypervisor.	12
APM Architecture.	13
About the APM Control Path.	15
About the APM Data Path.	16
About Multipathing and Device Mapper Automation.	16
Supported Distributions.	18
APM Requirements.	19
Operating Limits.	20

Chapter 2 Install APM

Prepare to Install APM.	21
Supported SAN Protocols.	22
Supported FC SAN Hardware Components.	22
Supported iSCSI Software and Hardware.	24
Pre-Configure the SAN Host for Pillar Axiom Integration.	29
About APM Software Installation.	31
Download the APM Software.	31
Install (or Update) the APM Software.	32
Configure SAN Host Access to the Pillar Axiom LUNs.	35
Partition and Format the LUN Disk (Optional).	38
About Load Balancing Configurations.	40
About Path Selection.	41
Remove the APM Software (Optional).	42

Chapter 3 APM Release Notes

New in this Release.	43
Known APM Issues.	44

Known Pillar Axiom Issues.	45
Known Operating System Issues.	47
Access to LUNs.	47
[13746418] SAN Dynamic Reconfiguration.	47
[13758951] Multipath Failed Path Errors.	50
[13757647] Unable to Update Path States.	50
[13764910] Non-Optimized Access Events After Slammer CU Failover.	51
[13761188] Not All iSCSI Paths Listed.	51
[13762121] Consult Oracle for Boot from SAN.	51
[13762235] Too Many Paths Cause Device Mapper to Hang.	52
[13762263] Oracle Cluster NOA Messages.	52
[13762246, 13765445] Call Trace During Path Failure.	52
[13762937, 13766418] LUNs Not Shown.	52
[13765459, 13766023] Host dom0 Reboots After Path Failure.	53
[13959093] LUNs Not Seen.	53
Resolved APM Issues.	54
Additional Notes.	55
Mount iSCSI Filesystems.	55
Index.	56



List of Figures

Figure 1 APM interaction with a Pillar Axiom server. 14

Figure 2 Example host ports before APM installation. 35

Figure 3 Example host ports after APM installation. 35



List of Tables

Table 1 Typography to mark certain content. 8

Table 2 Oracle resources. 9

Table 3 APM 3.2 for Oracle VM Server 2.2 features. 10

Table 4 Line and color key for APM interaction diagram. 13

Table 5 APM operating limits. 20

Table 6 Supported QLogic HBAs. 23

Table 7 QLogic HBA settings 23

Table 8 Characteristics of SAN connections to host ports. 30

Table 9 Known Pillar Axiom issues. 45

Preface

Related Documentation

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

Typographical Conventions

Table 1 Typography to mark certain content

Convention	Meaning
<i>italics</i>	Within normal text, words in italics indicate one of the following: <ul style="list-style-type: none">• A reference to a book title• New terms and emphasized words• Command variables
monospace	Indicates one of the following, depending on the context: <ul style="list-style-type: none">• The name of a file or the path to the file• <i>Output</i> displayed by the system on the command line
monospace (bold)	<i>Input</i> provided by an administrator on the command line.
>	Indicates a menu item or a navigation path in a 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).

Table 1 Typography to mark certain content (continued)

Convention	Meaning
...	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.

Oracle Contacts

Table 2 Oracle resources

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

CHAPTER 1

Introduction to APM

APM Features

Pillar Axiom Path Manager (APM) is defined as:

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

APM performs the following primary functions:

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

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

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

Table 3 APM 3.2 for Oracle VM Server 2.2 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 3.2 for Oracle VM Server 2.2 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 Oracle Customer Support. The information collected from each APM host includes:</p> <ul style="list-style-type: none"> • Logs from the APM components • Configuration and status information from the operating system • System and error logs from the operating system <p>No customer data is transmitted.</p>
Support for FC connections to FC Slammers	Makes connections to Pillar Axiom storage arrays over a 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 an IP network infrastructure.</p> <p>Note: iSCSI connections to FC Slammers require iSCSI-to-FC routers.</p>
Support for virtualization	APM supports the Xen hypervisor included in the Oracle VM Server distribution.
Support for load balancing	Both static and round-robin load balancing options are supported.

About APM and Xen Hypervisor

Pillar Axiom Path Manager (APM) supports the Xen hypervisor, an open source virtual machine monitor included with your operating system distribution.

You can install APM on an Oracle VM Server host configured as a Xen hypervisor host (dom0). Guest operating systems can access multipathed Pillar Axiom LUNs that are mapped to the dom0 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 a version of APM that supports iSCSI is available for the guest operating system, 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.

Refer to the [Oracle VM](http://www.oracle.com/us/technologies/virtualization/oraclevm/024974.htm) page (<http://www.oracle.com/us/technologies/virtualization/oraclevm/024974.htm>) for information on Oracle VM.

APM 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.

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

Table 4 Line and color key for APM interaction diagram




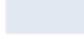

Graphic element	Description
———	Data path
-----	Control path
	Oracle-supplied hardware and software
	Non-Oracle hardware and software
	SAN host kernel space
	SAN host user space

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 path (all dashed lines)	14 HBA (FC) or NIC (iSCSI)
6 Pillar Axiom administrator	15 SCSI over FC (FC) or iSCSI over IP (iSCSI)
7 Pillar Axiom CLI or GUI	16 Data path (all solid lines)
8 Encrypted XML over TCP/IP	17 Pillar Axiom server
9 Network card	18 Brick storage pool

Related concepts

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

About the APM Control Path

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

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

- Get Fibre Channel (FC) and iSCSI port information from the HBA 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 storage system. In the Pillar Axiom Storage Services Manager, this description creates a definition that includes any FC ports in the host, and, if iSCSI is configured, the name of the iSCSI initiator. The graphical user interface (GUI) and command line interface (CLI) list the port World Wide Names (WWNs) of the FC ports in the host and the IP addresses that are used to make iSCSI connections to the Pillar Axiom system.

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

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

About the APM Data Path

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

The Linux multipath framework is responsible for the following:

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

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

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

Related references

- [APM Architecture](#)

About Multipathing and Device Mapper Automation

This release of Pillar Axiom Path Manager (APM) uses an updated Linux 2.6 `device-mapper` to provide multipathing services functionality.

APM consists of a daemon that monitors the state of the Pillar Axiom LUNs, and determines 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 a package containing an updated version of the `multipath-tools` component of `device-mapper`, along with additional bug fixes, as part of APM.

Important! You must use this Oracle version of `multipath-tools` instead of the one supplied as part of the operating system in order to use APM.

This release of APM also contains a package containing an updated version of the `device-mapper` that must be installed in addition to the updated version of `multipath-tools`.

In addition, 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 following Oracle-provided files:

- `/etc/init.d/multipathd`
- `/etc/init.d/axiompm`

Supported Distributions

Pillar Axiom Path Manager (APM) is supported on Oracle VM Server 2.2 platforms.

APM supports Oracle VM Server 2.2 distributions for the x86-32 and x86-64 (32-bit and 64-bit AMD and Intel platforms) architecture.

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

```
# uname -i
```

Verify that the output of this command is `i386`.

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:

```
2.6.18-128.2.1.4.37.el5xen
```

APM Requirements

Pillar Axiom systems presenting LUNs to Oracle VM Server 2.2 hosts using Pillar Axiom Path Manager (APM) 3.2 must be running release 4.5 or higher of the Pillar Axiom software.

Operating Limits

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

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

Table 5 APM operating limits

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

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

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

CHAPTER 2

Install APM

Prepare to Install APM

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

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

Supported adapters:

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

Related concepts

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

Related references

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

Related tasks

- [Install QLogic HBAs and Drivers](#)
- [Configure the iSCSI Initiator](#)
- [Start the iSCSI Software Initiator Service](#)
- [Pre-Configure the SAN Host for Pillar Axiom Integration](#)

Supported SAN Protocols

Pillar Axiom Path Manager (APM) supports Fibre Channel (FC) and iSCSI connections to the Pillar Axiom system.

Supported FC SAN Hardware Components

Pillar Axiom Path Manager (APM) supports a variety of Fibre Channel (FC) HBAs and SAN routers and switches.

Related references

- [Supported QLogic HBAs and Drivers](#)
- [Supported FC Switches](#)

Supported FC Topologies

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

Maximum Queue Depth

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

requests to each LUN on the Pillar Axiom system. Exceeding the maximum I/O queue depth can cause errors.

Supported QLogic HBAs and Drivers

Pillar Axiom Path Manager (APM) supports the following QLogic FC HBAs on x86-32 and x86-64 platforms:

Table 6 Supported QLogic HBAs

QLA2310F, QLA2310FL	QLE2460
QLA2340, QLA2340L	QLE2462
QLA2342, QLA2342L	QLE2464
QLA2344	QLE2560
QLA2460	QLE2562
QLA2462	QLE2564

Related references

- [Recommended Settings for QLogic HBAs](#)

Related tasks

- [Install QLogic HBAs and Drivers](#)

Recommended Settings for QLogic HBAs

We recommend the following settings for QLogic HBAs:

Table 7 QLogic HBA settings

HBA setting	Recommended value
Link Down Timeout	30
Execution Throttle	64
LUNs per Target	256
Port Down Retry Count	30

Table 7 QLogic HBA settings (continued)

HBA setting	Recommended value
Max Queue Depth (ql2xmaxqdepth)	64

Install QLogic HBAs and Drivers

To ensure compatibility between APM and your QLogic HBAs, be sure to follow the instructions in this section for installing the supported QLogic HBAs and drivers.

- 1 Install QLogic HBAs according to the instructions at the [QLogic support download page](http://driverdownloads.qlogic.com) (currently <http://driverdownloads.qlogic.com>).
- 2 Use the driver that ships with the Oracle VM Server 2.2 operating system.

Supported FC Switches

For a list of supported Fibre Channel (FC) switches, choose one of the following options:

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

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

Supported iSCSI Software and Hardware

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

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

Note: APM does not support iSCSI HBAs.

About Network Configuration for iSCSI Multipathing

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

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

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

About Subnet Configuration

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

About NIC Bonding

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

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

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

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

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

Configure the iSCSI Initiator

To support iSCSI connections, Pillar Axiom Path Manager (APM) requires the iSCSI software initiator included in your operating system 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 might include setting up Challenge Handshake Authentication Protocol (CHAP) on your system.
- 2 Configure iSCSI accordingly.
Example:
If you are using CHAP to communicate with the Slammer iSCSI ports, you may need to modify the `/etc/iscsi/iscsid.conf` file accordingly.
- 3 Start the iSCSI initiator.
Run the following command:

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

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

5 Log in to the Slammer iSCSI target ports.

Example:

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

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

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

Run the following command:

```
# service iscsi restart
```

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

Related tasks

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

Start the iSCSI Software Initiator Service

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

1 Start the `iscsi` service.

Run the following command:

```
# service iscsi start
```

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

Run the following commands:

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

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

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

Run the following command:

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

Result:

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

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

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

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

Supported iSCSI-to-FC Routers

iSCSI-to-Fibre Channel (FC) routing features enable a host to use 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:

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

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

Related references

- [Oracle 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 ports on the Slammer storage controller as iSCSI targets.

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

- 2 Present the iSCSI hosts as virtual 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, prepare your 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 iSCSI ports
- Supported HBA 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
- 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 HBA 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 AxiomONE 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 8 Characteristics of SAN connections to host ports

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

Related concepts

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

Related tasks

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

About APM Software Installation

After you prepare your SAN for Pillar Axiom Path Manager (APM), download and install the APM software. After that, configure APM and the Pillar Axiom software so that they can work together.

The APM installation requires that you download the APM package, the Multipath Tools package, and the Device Mapper package from My Oracle Support. Then you need to install all three packages on your system.

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

Removal of the APM software from your SAN host is a separate task.

Related tasks

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

Download the APM Software

Download the Pillar Axiom Path Manager (APM) software from My Oracle Support.

Prerequisite:

You need to be registered with My Oracle Support to download the APM software.

The APM software consists of three packages that you will need to download:

- APM: Contains the APM daemon
- Multipath Tools: Contains an updated version of the `multipath-tools` component of `device-mapper`
- Device Mapper: Contains an updated version of the `device-mapper`

The Multipath Tools package also contains bug fixes. When using APM, the Oracle version of the `multipath-tools` component must be used in place of the component supplied with the Linux operating system.

- 1 Sign in to [My Oracle Support](http://support.oracle.com) (<http://support.oracle.com>).
- 2 In the My Oracle Support dashboard, click **Patches & Updates**.

- 3 In the Patch Search frame, select **Product or Family (Advanced)**.
- 4 From the **Product is** list box, choose **Oracle Axiom Product Family**.
- 5 From the **Release is** list box, select the APM version.
- 6 From the **Platform is** list box, select the operating system and hardware platform.
- 7 Click **Search**.

Result:

Search results are displayed in the Patch Search Results window. If no results are returned, refine your search criteria and try again.

- 8 Click the name of the APM archive in the File Download dialog to begin the download.
- 9 Save the APM archive to your SAN host.
- 10 Extract the contents of the software bundle archive to your SAN host.

The archive contains APM, Multipath Tools, and Device Mapper software installation packages for all supported hardware platforms, as well as documentation, for the specified APM. Extract the APM, Multipath Tools, and Device Mapper packages for your hardware platform and (optionally) the documentation.

Install (or Update) the APM Software

After you download the 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 mentioned earlier in this document.

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 7 asks you to merge any changes you made to the original configuration file into the new file.

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

- 1 If you previously configured `multipath-tools` on your system, save a copy of the current `/etc/multipath.conf` file in a different directory.

Saving a copy allows you to merge your customizations into the new default file that will be created.

- 2 Remove any previous versions of APM on your system using the following command:

```
# rpm -e axiompmp
```

- 3 Remove any previously installed versions of `multipath-tools`.

The name of the `multipath-tools` Red Hat Package Manager (RPM) component varies (depending on its source). The name usually begins with one of the following strings:

```
multipath-tools
```

```
device-mapper-multipath
```

Example:

Use one of the following commands, depending on the name of the `multipath-tools` RPM:

```
# rpm -e multipath-tools
```

or

```
# rpm -e --nodeps device-mapper-multipath
```

Important! Other packages associated with `device-mapper` that begin with the string `device-mapper` may exist on the SAN host. Do not remove these packages. Remove only those packages that begin with the string `multipath-tools` or `device-mapper-multipath`.

- 4 Install the Device Mapper package to upgrade the `device-mapper` using the following command:

```
# rpm -Uvh --nodeps device-mapper-1.02.67-2.el5.i386.rpm
```

- 5 Install the Multipath Tools package using the following command:

```
# rpm -ivh multipath-tools-version.rpm
```

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

- 6 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.

- 7 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.

- 8 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 FC SAN Hardware Components](#)

Related tasks

- [Pre-Configure the SAN Host for Pillar Axiom Integration](#)
- [Configure SAN Host Access to the Pillar Axiom LUNs](#)

Configure SAN Host Access to the Pillar Axiom LUNs

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

Prerequisites:

- Pre-configure the SAN host for Pillar Axiom integration.
 - Install the Pillar Axiom Path Manager 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.
- 4 As needed, set up mappings of the LUNs to the new host entries.
- 5 Make any changes visible to the SAN host by following the recommendations for SAN dynamic reconfiguration.
- 6 (Optional) Run the following commands to list the multipath devices:

```
# /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

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

Related tasks

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

Partition and Format the LUN Disk (Optional)

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

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

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

Example:

In the output below, you could identify either the `/dev/sdd` 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 the load in different ways:

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

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

About Path Selection

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

Paths to a LUN can 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 might remain lower than before because 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 the APM 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 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 axiomp
```

Sometimes, the uninstall process may take longer than expected, and the process may appear to be hanging. Allow the process to run for at least an hour before attempting to restart the process.

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). You also may need to reconfigure the host to access the LUNs through new device names.

CHAPTER 3

APM Release Notes

New in this Release

Pillar Axiom Path Manager (APM) 3.2 supports Oracle VM Server 2.2 for x86.

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 9 Known Pillar Axiom issues

Issue	Workaround or planned fix
<p>[13759030] If an iSCSI initiator is connected to a port on a Slammer, and that Slammer CU is powered off, the LUN and Host GUI pages continue to show the connection status for the iSCSI initiator as <code>Connected</code>.</p>	<p>This issue is fixed in release 5.0 of the Pillar Axiom software.</p>
<p>[13759805] If more than 256 SAN LUNs are configured on a Pillar Axiom system, the Pillar Axiom Storage Services Manager may send invalid messages to the APM daemon running on SAN host systems. When this happens, the control path connection between APM and the Pillar 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.</p>	<p>This issue is fixed in release 5.0 of the Pillar Axiom software.</p>
<p>[13764561] The Pillar Axiom Storage Services Manager (GUI) sometimes continues to display host paths that are no longer valid after APM stops. 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>[13764609, 13762326] While the system is recovering from temporary</p>	<p>This issue is fixed in release 4.5.1 and 5.2.1 of the Pillar Axiom software.</p>

Table 9 Known Pillar Axiom issues (continued)

Issue	Workaround or planned fix
use of non-optimized paths to the alternate CU on a Slammer, you may notice a decrease in I/O performance between a LUN and a Slammer CU.	<p>If the Pillar Axiom system is running a release earlier than 4.5.1 or 5.2.1, you can take the following actions:</p> <ol style="list-style-type: none"><li data-bbox="966 436 1414 506">1 Verify your SAN connections to restore the optimized paths.<li data-bbox="966 527 1425 659">2 Follow the instructions in the <i>Pillar Axiom Administrator's Guide</i> to rehome the LUN to the alternate CU on that Slammer.<li data-bbox="966 674 1398 743">3 Rehome the LUN again to the original CU on that Slammer. <p>Contact Oracle Customer Support for assistance.</p>

Known Operating System Issues

The following operating system issues can have an impact on running Pillar Axiom Path Manager (APM) on Linux and Oracle VM Server 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 \(Optional\)](#)

[13746418] SAN Dynamic Reconfiguration

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

Follow the instructions for dynamic reconfiguration in the [Online Storage Reconfiguration Guide](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Online_Storage_Reconfiguration_Guide/) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Online_Storage_Reconfiguration_Guide/). If you continue to have problems with dynamic reconfiguration, the following steps will help you handle exceptions.

Linux requires a series of steps, including a potential host reboot, when a LUN is dynamically reconfigured. After Linux detects the new LUNs, you must restart

Pillar Axiom Path Manager (APM) to update the Pillar Axiom system with the new LUN status.

LUNs Added Dynamically

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

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

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

LUNs Deleted Dynamically

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

If a LUN is deleted, it may appear as either a `2000000000000` entry or as the original LUID with `Path down` messages. These entries may persist until the host is rebooted.

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

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

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

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

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

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

```
#!/bin/bash
# Must be run as root
```



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

Tip: The rescan script might require your interaction.

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

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

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.

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 Reconfiguration Guide](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Online_Storage_Reconfiguration_Guide/) (http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Online_Storage_Reconfiguration_Guide/) to force your device drivers to recognize the resized LUN, the device-mapper may recognize the resized LUN without a host reboot. The ability to use the resized LUN is a function of the host filesystem.

Clone LUNs Added or Deleted Dynamically

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

[13758951] Multipath Failed Path Errors

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

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

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

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

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

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

[13757647] Unable to Update Path States

After paths are lost and restored, it may be necessary to manually refresh `device-mapper` to update the host path state.

Because of a Linux defect, it may be necessary to run the following command to restore the path state after paths have been lost and restored on the host:

```
/sbin/multipathd reconfigure
```

Running the `/sbin/multipathd show maps topology` command to show the state of the LUNs as seen by the kernel fails to show the current path status on the Linux host. This prevents the Pillar Axiom Path Manager (APM) daemon from updating the LUN status.

If running the `multipathd reconfigure` command does not recover the condition, run the following command to restart the multipath daemon:

```
service multipathd restart
```

[13764910] Non-Optimized Access Events After Slammer CU 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.

[13761188] 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, rescan all iSCSI devices using the following command:

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

Restart the `multipathd` service using the following command:

```
service multipathd restart
```

[13762121] Consult Oracle for Boot from SAN

Consult Oracle Customer Support for specific instructions and details of configuring boot from SAN.

[13762235] Too Many Paths Cause Device Mapper to Hang

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

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

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

```
service multipathd restart
```

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

[13762263] Oracle Cluster NOA Messages

When a node in a cluster reboots, it may take up to six minutes for paths to switch to the optimized paths, causing transitional non-optimized access (NOA) messages to be generated.

This is a defect in the `device-mapper` code shipped by Oracle.

In Steady state, I/O runs in the optimized path as expected.

A service request SR 3-3696665562 has been filed with Oracle.

[13762246, 13765445] 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.

[13762937, 13766418] LUNs Not Shown

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

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

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

[13765459, 13766023] Host dom0 Reboots After Path Failure

After a path failure between the host and the Pillar Axiom system, the host may unexpectedly crash and reboot.

The path failure may be caused by a failure in the Fibre Channel (FC) or FC over Ethernet (FCoE) SAN, or by a failure in the Pillar Axiom system such as a Slammer control unit (CU) being restarted or powered off.

Contact Oracle Customer Support for a solution.

[13959093] LUNs Not Seen

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

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

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

Resolved APM Issues

There are no resolved APM issues for this release.

Additional Notes

Mount iSCSI Filesystems

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

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

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

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

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

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

- 3 Restart the system.

Result:

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

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

Index

A

- additional notes *55*
- additional resources *8*
- APM
 - control path *15*
 - how to
 - install software *32*
 - remove software *42*
 - update software *32*
 - interaction with Pillar Axiom *13*
 - pre-installation requirements *21*

APM software

- data path *16*
- definition *10*
- features *10*
- how to
 - download *31*
- known issues *44*
- operating limits *20*

C

- Cisco MDS switch
 - configure *29*
- configuration
 - iSCSI multipathing *25*
 - iSCSI-to-FC router *29*
 - Pillar Axiom LUNs *35*
 - SAN dynamic reconfiguration *47*
 - subnet *25*
- contact information *9*
- contacts, Oracle *9*
- conventions
 - typographical *8*
- customer support *9*

D

- device-mapper
 - function *16*
- documentation *9*
 - feedback *9*
- download
 - APM software *31*

- multipath-tools *31*

E

- education programs *9*

F

- FC
 - switches *24*
- features
 - APM *10*
 - new in this release *43*
- feedback, documentation *9*
- Fibre Channel
 - iSCSI-to-FC routers *28*
 - protocol *22*
 - supported HBAs *22*
 - supported topologies *22*
- format
 - LUN disk *38*

H

- HBAs
 - QLogic *23, 24*

I

- install
 - APM software *32*
 - multipath-tools *32*
- iSCSI
 - how to
 - configure Cisco MDS switch *29*
 - configure iSCSI initiator *26*
 - configure iSCSI-to-FC router *29*
 - start iSCSI initiator service *27*
 - NIC bonding *25*
 - supported hardware *24*
 - supported routers *28*
 - supported software *24*
 - supported switches *28*
- iSCSI initiator
 - start service *27*
- iSCSI multipathing

- subnets *25*
- iSCSI-to-FC router
 - configure *29*
- issues
 - additional notes *55*
 - APM software *44*
 - operating system *47*
 - Pillar Axiom system *45*
- L**
- load balancing
 - round-robin *40*
 - static *40*
- LUN access
 - multi-protocol *41*
 - precedence *41*
 - preferred paths *41*
- LUN configuration
 - SAN hosts *35*
- LUN disk
 - how to
 - format *38*
 - partition *38*
- M**
- multi-protocol access
 - definition *41*
 - precedence *41*
 - preferred paths *41*
- multipath-tools
 - function *16*
 - how to
 - download software *31*
 - install software *32*
 - remove software *42*
 - update software *32*
 - pre-installation requirements *21*
- N**
- new features *43*
- NIC bonding *25*
- non-optimized path *41*
- O**
- online help *9*
- operating system
 - known issues *47*
- optimized path *41*
- Oracle Technical Network (OTN) *9*
- P**
- partition
 - LUN disk *38*
- Pillar Axiom
 - configure for LUN access *35*
 - interaction with APM *13*

- Pillar Axiom system
 - known issues *45*
 - requirements *19*
- pre-configure SAN hosts *29*
- pre-installation requirements *21*
- product support *9*
- Q**
- QLogic HBAs
 - install *24*
 - settings *23*
 - supported *23*
 - supported drivers *23*
- queue depth
 - maximum for SAN host HBAs *22*

- R**
- related documentation *8*
- release notes
 - additional notes *55*
 - APM issues *44*
 - operating system issues *47*
 - Pillar Axiom issues *45*
- remove
 - APM software *42*
 - multipath-tools *42*
- requirements
 - Pillar Axiom system *19*
- round-robin load balancing *40*
- routers
 - iSCSI-to-FC *28*

- S**
- sales information *9*
- SAN hosts
 - LUN configuration *35*
 - maximum HBA queue depth *22*
 - pre-configure *29*
- SAN protocols *22*
- static load balancing *40*
- subnets
 - iSCSI multipathing *25*
- Support portal *9*
- supported
 - architectures *18*
 - Fibre Channel topologies *22*
 - hardware platforms *18*
 - iSCSI hardware *24*
 - iSCSI software *24*
 - Linux kernels *18*
- switches
 - Cisco MDS *29*
 - FC *24*

iSCSI *28*

T

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

U

uninstall
 APM software *42*
 multipath-tools *42*
update
 APM software *32*
 multipath-tools *32*

V

virtualization
 Xen hypervisor *12*

X

Xen hypervisor *12*