Pillar Axiom 500 and 600

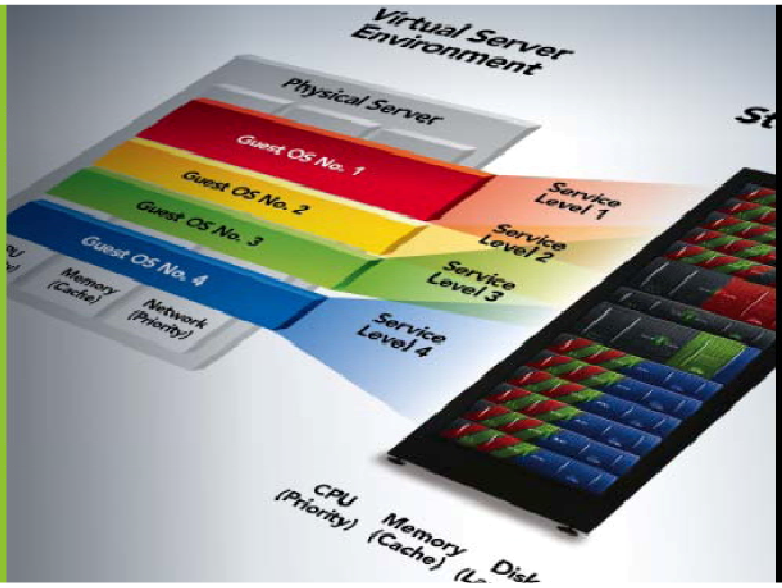

Customer Release Notes

For Release 5.2

ORACLE

PILLAR AXIOM

Pillar Axiom 500 **Pillar Axiom 600**	**Document Number: 4420-00026-4300** **Document Title: Customer Release Notes**

Revision History

Rev Description	Rev Date	Effective Date
Release 5.2	2011-09-30	2011-10-01 [Pillar Axiom 500 and 600 SAN only]
Release 5.0 GA Rev 1.1	2011-06-21	2011-06-21 [Pillar Axiom 500 and 600 SAN only]
Release 5.0 GA	2011-05-19	2011-06-13 [Pillar Axiom 500 and 600 SAN only]
Release 4.3.1 GA Rev 1.1	2011-02-08	2011-02-11
Release 4.3.1 GA	2011-01-03	2011-01-19
Release 4.3.0 Beta	2010-12-14	2010-12-28
Release 4.2	2010-06-21	2010-07-22
Release 4.1	2010-03-30	2010-04-27
Release 4.0	2009-11-20	2009-12-21
Release 3.4 GA	2009-07-22	2009-08-28 [Pillar Axiom 600 only]
Release 3.3 GA	2009-03-20	2009-05-01
Release 3.2 GA	2008-10-20	2008-10-27
Release 3.2 Beta	2008-10-10	2008-10-20
Release 03.01.00	2008-06-06	2008-06-18
Release 03.00.00	2008-03-12	2008-03-14

1 Terms and Conditions of Use

All systems are subject to the terms and conditions of the software licensing agreements and relevant copyright, patent, and trademark laws. Refer to those documents for more information.

2 Purpose

This document describes new features, capacities, configuration requirements, operating constraints, known issues and their workarounds, and other items for release 5.2 of Oracle's Pillar Axiom 500 and 600 storage systems. The document covers hardware, firmware, software, cabling, and documentation. The information provided is accurate at the time of printing. Newer information may be available from your Pillar Axiom authorized representative.

3 Product Release Information [a]

Release 5.2 is a special maintenance release of the Pillar Axiom software for the Pillar Axiom 500 and Pillar Axiom 600 SAN-only platforms.

3.1 System Enhancements

The overall system functions provided in this update are similar to those provided in release 5.0. In addition, this release provides improved system efficiency and robustness, as well as including a rollup of all defects fixed in customer patches up to and including release 5.2.

For the list of defects that this release resolves, see Table 8 beginning on page 26.

3.2 Changes to Licensing

All licensing requirements in Pillar Axiom systems, with the exception of the secure shell (SSH) license, have been removed. For more information on licensing, refer to Section 3.4, Licensing Optional Premium Features.

The SSH license can still be used by Customer Support Center personnel.

3.3 Changes to How the Pillar Axiom Software Operates

3.3.1 SAN Replication

The AxiomONE Replication for SAN product is supported in 4.x releases only. It is not available for or supported in release 5.0 and higher. The Pillar Axiom MaxRep Replication for SAN product, which supports asynchronous and synchronous replication, is available as of release 5.0.

The current SAN replication product, which is named Pillar Axiom MaxRep Replication for SAN, is *not* compatible with AxiomONE Replication for SAN. All customers who use SAN replication must utilize Professional Services to establish a new installation or migrate from an existing replication solution.

3.3.2 Capacity Utilization Across Brick Types (Storage Classes)

In software releases prior to release 4.1, for a Pillar Axiom system containing a mix of SATA and Fibre Channel (FC) Bricks, individual logical volumes could utilize space from both the SATA and FC storage pools. As of release 4.1 and later releases, administrators cannot *create* volumes that span Brick types (called *Storage Classes* as of release 4.1). Legacy volumes created prior to release 4.1 that span Storage Classes, however, can continue to exist in the system after updating to release 5.2.

Important! Beginning with release 4.1, administrators cannot grow these legacy volumes (or in-fill thinly-provisioned volumes) that span multiple Storage Classes unless available capacity exists for doing so in the Storage Class on which the volume was originally created.

[a] Throughout this document, all references to release 2.x apply to the Pillar Axiom 500 system.

Also, beginning with release 4.1, if legacy volumes that span the SATA and FC Storage Classes exist, the available capacity shown in the management interfaces (the GUI and the CLI) will be reported differently. These interfaces now report the capacity that is available by Storage Class, which is more accurate.

As of release 5.0, to accommodate a legacy logical volume that spans the SATA and FC Storage Classes, you have to have sufficient SATA space within the Storage Domain in which the volume resides.

3.3.3 Reported Free Capacity

Beginning with release 5.0, the free, available, and total system capacities reported by the system are higher by either 50 GB or 100 GB. The reason for this change is as follows:

The Pillar Axiom system allocates 50 GB of physical capacity in the SATA and FC (but not SSD) Storage Classes as an "in-fill reserve". The system reserves this physical capacity to help prevent inadvertent exhaustion of system physical capacity when LUNs have been created with Thin Provisioning. This capacity would be used when physical capacity needs to be assigned to a thinly provisioned volume, and all other physical capacity in that Storage Class has been consumed.

In releases prior to 5.0, the size of this reserve capacity *was not included* in calculations of system-wide capacities. As of release 5.0, the size of this reserve capacity *is included* in these calculations. This change means that free, available, and total system capacities as of release 5.0 are either 50 GB or 100 GB larger in physical capacity than the corresponding figures in release 4.x, depending on whether the Pillar Axiom system had one or both SATA and FC Storage Classes present prior to the software update.

3.3.4 Method of Reporting Volume Sizes

As of release 5.0, what had been reported as *Current Capacity* and *Maximum Capacity* has been changed to *Allocated Logical Capacity* and *Addressable Logical Capacity*, respectively. These two capacities are defined in the following way:

- **Allocated logical capacity.** The amount of storage physically allocated for the logical volume. This value can change if the logical volume is migrated or copied with a different Quality of Service setting or if Thin Provisioning in-fill occurs. This value is the actual amount of storage allocated, minus RAID overhead.
- **Addressable logical capacity.** The size of the logical volume as visible to a client. This capacity defines the LBA range visible to a client and is guaranteed to remain constant across data migrations or volume copies that do not involve a volume resize operation.

3.3.5 Network TCP Port Configuration

After updating a Pillar Axiom system to release 5.2, administrators then access the updated system from a client workstation to download various utilities, including the new Java-based graphical user interface (GUI), Pillar Axiom Storage Services Manager. (See Section 3.5.3.2.)

To download the utilities, the client uses HTTP to access the Pillar Axiom system. If an internal firewall is in place, you need to ensure that the TCP port 26008 in the firewall is open to allow TCP/IP traffic between the client host and the Pilot.

Tip: This TCP port must also be open to enable the Pillar Axiom CLI client to gain access through the firewall to the Pillar Axiom Pilot.

3.4 Licensing Optional Premium Features

All features on the Pillar Axiom 600 storage system are enabled out of the factory. Administrators should ensure they are in compliance with their End User License Agreements and have purchased the necessary licenses for Optional Premium features.

The following features are currently licensed on the Pillar Axiom 600 storage system:

- Pillar Axiom Storage Domains - System Perpetual
- Pillar Axiom Copy Services Bundle - System Perpetual

The following features are currently licensed on the Pillar Axiom Replication Engine:

- Pillar Axiom MaxRep Asynchronous Replication - Terabyte Perpetual
- Pillar Axiom MaxRep Asynchronous Replication with Application Protection - Terabyte Perpetual
- Pillar Axiom MaxRep Synchronous Replication - Terabyte Perpetual
- Pillar Axiom MaxRep Synchronous Replication with Application Protection - Terabyte Perpetual

For additional information, please contact Pillar Data Systems at 1-877-4PILLAR or go to www.pillardata.com.

3.5 Pillar Axiom Software Update

The 5.2 release may be installed by means of the GUI software update process.

Important! To update a Pillar Axiom system to release 5.2, the system must be running release 4.3.15 (or higher) software. Updating a Pillar Axiom system from release 4.3.15 (or higher) to release 5.2 is *disruptive* to the data path. Also, the update to release 5.2 is not reversible.

Refer to the *Pillar Axiom Customer Release Notes for Release 4.3* for information on updating earlier systems to release 4.3.15.

Important! All customers who use Pillar Axiom MaxRep for SAN must contact Professional Services to establish a new installation or migrate from an existing replication solution. (For more information, see Section 3.3.1.)

If you prefer, you can have Pillar Professional Services update your system software. For more information, refer to Table 3 Contact information.

3.5.1 Software Update Packages

When updating a Pillar Axiom system from release 4.3.15 to release 5.2, you use the following traditional types of software package:

```
Ax500_FN_050200-001300.tgz
Ax600_FN_050200-001300.tgz
```

Those files are hardware specific packages, for which the first five characters of the file name indicate the type of hardware on which the package is to be installed. These packages are compressed tar files.

After the system has been updated to release 5.2, hardware specific packages (tgz files) will no longer be used for updates to higher software releases. Instead, to update release 5.x software, you will use the rpm style package format, which applies to both the Pillar Axiom 500 and 600 platforms. The rpm style package format looks like the following:

```
AxiomONE-SW-050200-001300.i386.rpm
```

3.5.2 The Release 5.2 Update Process

The Pillar Axiom system must be in a green and Normal status and have no outstanding Administrator Actions (for 4.x systems) or system alerts (for 5.x systems) *before* you begin the process of updating your system to release 5.2.

Note: If the currently installed system software is an R5 release below 05.00.05, Pillar recommends that the update be installed with the Restart System option.

Important! After updating the system software to release 5.2, you cannot reverse the update and downgrade to an earlier release level without explicit action by the Pillar World Wide Customer Support Center.

3.5.3 After the Update

After a successful update, all system components report Normal and user data is available. At this point, several additional actions must be performed to make the Pillar Axiom system fully operational.

3.5.3.1 Client TCP Port Configuration

After updating a Pillar Axiom system to release 5.2, administrators need to access the system from a client workstation and download various utilities, such as the new Java-based graphical user interface (GUI), Pillar Axiom Storage Services Manager. (See Section 3.5.3.2.)

To download the GUI and other utilities, the client uses HTTP to access the Pillar Axiom system. If an internal firewall is in place, you need to ensure that TCP port 26008 is open to allow the TCP/IP traffic to occur between the Pillar Axiom system and the client.

This TCP port must also be open so that the encrypted traffic between the Pillar Axiom CLI clients and the Pillar Axiom system can pass through the firewall.

3.5.3.2 Install the Pillar Axiom Storage Services Manager

Tip: You can view system status and recent events using an HTTP connection to the Pillar Axiom system, which can be done from a mobile device if desired. Simply enter the IP address or the name of the system and browse for the appropriate link.

Using this HTTP connection, you cannot modify the configuration of the Pillar Axiom system. To configure and otherwise manage the system, you need to download, install, and run the Pillar Axiom Storage Services Manager GUI, which is a Java-based client application.

For instructions on downloading, installing, and running the Pillar Axiom Storage Services Manager GUI, refer to the *Pillar Axiom Administrator's Guide*. You can download this guide from your Pillar Axiom system as follows:

1. Point your browser to http://*system-name-IP*/documentation.php, where *system-name-IP* is the name or the public IP address of your Pillar Axiom system.
2. Click the **Technical Documentation** link.
3. Click the **Administrator's Guide** link.

Note: If you have an existing installation of the Pillar Axiom Storage Services Manager (GUI) client and attempt to log in to the Pillar Axiom system after a software update, you will be prompted to update your client. Please do the update, as described in the *Pillar Axiom Administrator's Guide*.

3.5.3.3 Pillar Axiom GUI Supported Platforms

The Pillar Axiom Storage Services Manager GUI is supported on the following platforms:

- Windows XP
- Windows Vista
- Windows 7
- Windows Server 2003
- Windows Server 2008
- Linux Fedora Core
- Ubuntu

Linux and Windows platforms require Java version 1.6.0 or higher. MacIntosh platforms require Java version 1.6.0_24.

Note: For Windows, the MSI installer does not require that Java be installed.

3.5.3.4 Pillar Axiom CLI Supported Platforms

For release 5.2, the Pillar Axiom CLI is supported on the following platforms

- Citrix 5.6 XenServer x86_64
- Mac OS X 10.6 x86_64
- Open SuSE 11.3
- Red Hat/CentOS/Oracle Linux, versions 4 and 5
- SLES 11 x86
- Solaris 9 on Sparc
- Solaris 10 on Sparc
- Solaris 10 on x86
- Windows 2003
- Windows 2008

- Windows 7
- Windows Vista
- Windows XP

3.5.4 Software Versions for This Release

Software for this release includes the versions listed in the following table. After updating your system, check your software versions in the GUI by going to **Support > Software Modules**.

Table 1 Pillar Axiom Software versions

Software module		Pillar Axiom model	Software version
Pilot OS		All	05.02.00
Pilot Software		All	05.02.00
Slammer PROM		500	03.00.00
		600	05.00.00
Slammer Software		All	05.02.00
Serial ATA (SATA) Brick Firmware	SATA RAID controller, Version 1	All	07.20.19
	SATA V2 RAID controller, Version 2 (includes SSD)	All	00.20.19
Fibre Channel (FC) Brick Firmware	FC RAID controller, Version 1	All	01.20.19
	FC V2 RAID controller, Version 2		02.20.19
Brick Disk Drive Firmware		*Version depends on drive capacity and whether it is the spare or an array drive.*	

Note: If your Pillar Axiom system does not contain a particular Brick type (for example, Fibre Channel, SATA, SATA V2, or SSD), the GUI does not display the firmware entry for that Brick type.

3.6 Pillar Axiom Path Manager Software

The Pillar Axiom Path Manager (APM) software provides the following features:

- Automatic data path failover
- Automatic recognition of SAN hosts in the Pillar Axiom Storage Services Manager
- Management of the APM driver from within the Pillar Axiom Storage Services Manager

The current release of APM for SAN hosts varies by platform, as shown in Table 2.

Table 2 APM support

Operating System	OS Release	APM Release	Platforms
AIX	5.3 6.1 7.1	3.1	All platforms
Citrix XenServer	5.6	3.0	All platforms

Community Enterprise Operating System (CentOS)	4.8	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86
	5.6	3.4	32-bit x86, 64-bit x86
HP-UX	11i v3	2.0	All 64-bit platforms
Oracle Linux (OL)	4 8	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86
	5.6	3.4	32-bit x86, 64-bit x86
Oracle Solaris	9	2.1	SPARC
	10	2.0	SPARC, 64-bit x86
Oracle VM Server for x86	2.2	3.1	All platforms
Red Hat Enterprise Linux (RHEL)	4.8	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86, POWER
	5.6	3.4	32-bit x86, 64-bit x86
SUSE Linux Enterprise Server (SLES)	10 SP2	3.0	32-bit x86, 64-bit x86
	11 SP1	3.0	32-bit x86, 64-bit x86
Windows	Server 2003 Server 2003 R2 Server 2008 Server 2008 R2	3.3	All platforms

Tip: For optimal performance with the 5.2 release of the Pillar Axiom software, Pillar recommends that APM clients use round-robin access from the host. Doing so ensures that all ports and Pillar Axiom CPUs are fully utilized.

Note: Unless otherwise explicitly stated in your APM Release Notes, there are no co-requisite relationships between the Pillar Axiom Path Manager and the Pillar Axiom software versions.

For release information, refer to the *Pillar Axiom Path Manager Installation Guide and Release Notes* for your platform. For the latest information on supported platforms and hardware, see the *Pillar Axiom Support and Interoperability Guide* or ask your Pillar Data Systems representative.

4 Support

Various levels of customer service are provided on a contract basis for Oracle's Pillar Axiom storage systems. If you have purchased a service contract from Oracle or Pillar Data Systems, authorized support personnel will perform support and repair according to the terms and conditions of that agreement.

Table 3 Contact information

For help with…	Contact…
Technical Support	U.S. and Canada: 877-4PILLAR (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. Email: support@pillardata.com Web: Customer support portal (http://support.pillardata.com/)
• Implementation assistance • System information • Enhancement requests	sales@pillardata.com USA: 1-877-4PILLAR (1-877-474-5527) Request Sales at the prompt. International: +1 408 503 4200
Documentation improvements and resources	docs@pillardata.com www.pillardata.com/techdocs Log in with your username and password.

4.1 Supported Hardware Components in a Pillar Axiom System

Pillar Data Systems supports only Pillar-supplied parts for Pillar Axiom systems. Hardware that does not conform to Pillar specifications or is not a Pillar-supplied part voids the warranty and might compromise data integrity.

4.2 Access to Pillar Axiom Systems

You manage a Pillar Axiom system by means of the standard user interfaces:

- The Pillar Axiom Storage Services Manager (GUI)
- The Pillar Axiom Command Line Interface (CLI)

Remote access by any other means (ssh, telnet, ftp, and others) is not supported and voids the warranty for your Pillar Axiom system. Furthermore, remote access may also compromise integrity of data that is stored on the system.

4.3 Download Software or Firmware Updates

To download software or firmware updates:

1. Point your browser to the Customer support portal. (http://support.pillardata.com)
2. Enter your registered username and password.

   Tip: After you log in, the portal displays your name and your company name in the upper right corner.
3. In the menu bar, select **My Downloads** > **My Software Releases**.
4. If you have more than one Pillar Axiom system, click the **System** drop-down list and select the serial number of the system for which you want the software update.
5. Under Folders for All Software, navigate to the software type you want and click the release you want to download.

   If you are a value-added reseller and the software release is not displayed for your system, it could be because you logged in as the wrong customer. Otherwise, contact the Pillar World Wide Customer Support Center for help.
6. In the **Available Software** content pane, click the title of the software package that you want to download.
7. In the **Software Download Information** content pane, review the details of the download package to verify your selection.

   Note: Check the file size of the download and be sure your local system has sufficient space.
8. To download the software package, click the **Download** link.

   The system displays a dialog box to inform you that the software package is being read. This operation can take a few minutes, depending on the size of the package.

   Tip: If your browser window displays an information bar that states that the download is blocked, click the appropriate options to allow the download to proceed.
9. To begin the download, click **Save**.

   Browse to the location on your local system where you want to save the software update package.

   If your client system offers the option to rename the file or to open the file with WinZip or another utility, select No and save the file with the original name.
10. To save the software package on your local system, click **Save**.

    After you successfully download the software update package, in the original dialog box, click the **Close When Finished** link.

4.4 Configuration Documentation

For information on the connectivity and interoperability of Pillar Axiom systems with various third-party software and hardware, see your Pillar Account Representative.

For detailed configuration guidelines in an iSCSI environment, see the *Pillar Axiom iSCSI Integration Guide for SAN Systems*.

For information regarding the primary features of a Pillar Axiom system and how to configure them:

- Navigate through the Pillar Axiom Storage Services Manager GUI.
- Read the *Pillar Axiom Administrator's Guide* PDF.
- Read the online help in the Pillar Axiom Storage Services Manager GUI.

The above documents can be obtained in any one of the following ways:

- In the GUI, navigate to **Support > Documents.**
- 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.
- Log in to the Customer support portal and click **Documents** in the left navigation pane. (http://support.pillardata.com/)
- Insert the Technical Documentation CD-ROM (which came with your Pillar Axiom system) into the CD player and open the `DocMenu` PDF.

5 Pillar Axiom System Limits

This version of the Pillar Axiom system operates within the supported limits listed below.

Important! Use care when operating a system that has been configured to run at or near the system operating limits. The system may exhibit anomalies when all limits are exercised concurrently. Also, the time to start Pillar Axiom systems from a powered-off or shutdown state and the responsiveness of the GUI are extended under the following conditions:

- You configure a system near one or more of its limits.
- You increase the number of customer-defined system objects, such as LUNs, clones, and so on.

Consult with Pillar Professional Services to plan your Pillar Axiom system configuration prior to actual installation and configuration.

5.1 Pillar Axiom System Operating Limits

For detailed information on system limits, refer to the online help or to the *Pillar Axiom Administrator's Guide* PDF file (search for *Ranges for Field Definitions*).

Table 4 Operating limits for all Pillar Axiom systems

Item	Description and range
Volume groups	Minimum = 1 Maximum = • 5000 total, out to five levels • 100 subgroups for a given volume group
Repository VLUNs[b]	Maximum = 1024
VLUNs	Maximum = • 8191 for each Slammer • 8191 for each system
Storage Domains	Maximum: 64 for each system
Number of Bricks in a Storage Domain	Minimum: • Serial ATA (SATA) or solid state drives (SSD) Bricks: 1 • Fibre Channel (FC) Bricks: 2 Maximum: • SATA Bricks: 64 • FC or SSD Bricks: 32

[b] A repository VLUN is associated with a logical volume and holds metadata for clones of that volume. A volume has at most one repository VLUN associated with it.

Table 5 Operating limits for Pillar Axiom SAN systems

Item	Description and range
SAN LUNs	Maximum = • 8191 visible for each system • 256 visible for each host • 8191 visible for each SAN Slammer
SAN LUN size	Minimum: 1 to 2 GB. The exact value depends on these factors • Brick type (Fibre Channel, SATA, or SSD) • RAID geometry (RAID 5 or Distributed RAID) • Strip size (1 MB or normal) Maximum = system capacity (unlimited when using Thin Provisioning)
Volume Copies (full block snapshots)	Maximum = number of unallocated SAN LUNs, up to maximum SAN LUNs • 1024 for each system • 12 active for each LUN
Clone LUNs (partial block snapshots)	Maximum = number of unallocated SAN LUNs, up to maximum SAN LUNs
Clone repositories	Maximum = ¼ of the maximum number of SAN LUNs
iSCSI	Maximum = • 256 TCP connections for each iSCSI port • 256 iSCSI initiators for each iSCSI port • 2048 iSCSI initiators for each Slammer • 256 LUNs for each initiator • 32 persistent reservation registration keys for each LUN • 512 simultaneous commands for each iSCSI port
LUN mappings	Maximum = 512K for each Pillar Axiom system

5.2 Slammer and Brick Configuration Limits

The minimum and maximum configurations for Pillar Axiom 500 and 600 systems are summarized in the following table:

Table 6 Brick configuration limits

Number of Slammers	Minimum number of Bricks		Maximum number of Brick strings	Maximum number of Bricks
	Supported	Recommended		
1	1	3	4	32
2	4	5	8	64
3	6	8	8	64
4	8	8	16	64

Note: The maximum number of SSD Bricks depends on the number of Slammers:

- 1 Slammer = 8 Bricks
- 2 and 3 Slammers = 16 Bricks
- 4 Slammers = 32 Bricks

Note: The maximum number of Bricks in any string is 8.

For Fibre Channel (FC) Bricks:

- Version 1 FC Bricks are available as FC RAID Bricks and FC Expansion Bricks. Each Version 1 FC RAID Brick supports the attachment of a single Version 1 Expansion Brick.

  Note: Version 2 FC Bricks do not have an FC Expansion Bricks option.
- Pillar Axiom systems have a limit of 32 FC Bricks, regardless how many Slammers are in the system.
- A given Brick string can contain up to four FC Bricks (RAID or Expansion). Maximum number of FC Expansion Bricks in any string, however, is two.

For a complete list of the rules for configuring FC, SATA, and SSD Bricks, see the appropriate *Pillar Axiom SSF Cabling Reference* for your Pillar Axiom system.

6 System Requirements

6.1 Windows Requirements

6.1.1 Using Browsers on Windows XP Operating Systems

Management of a Pillar Axiom system is done over a secure connection using the Java based client application.

However, when logging into the Pillar Axiom Storage Services Manager using secure HTTP (HTTPS), for example to download the GUI, utilities, or documentation or to check system status or recent alerts, you might see warnings that the server certificate is not issued by a trusted authority. The server certificate is installed and signed by Pillar Data Systems during the manufacturing process.

If this Pillar certificate is not suitable for your security requirements, server certificates are available for purchase from a number of Certificate Authorities. Be sure any installed CA certificate is not password protected.

6.1.2 Memory Requirements

For the Pillar Axiom Storage Services Manager graphical user interface to operate correctly, ensure that the client host has a minimum of 2 GB memory.

6.2 Network Requirements

6.2.1 Pilot Network Requirements

The Pilot management controller requires:

- Two 100 BaseT ports for the public connection to the management network. For added redundancy, the two connections should be to separate switches. The Pillar Axiom system provides a standard Cat 5 RJ-45 jack on each Pilot control unit (CU) for this connection.
- Three IP addresses on the same subnet: one IP for each physical interface and one shared IP.

Note: VLAN tagging is not supported on the management interfaces.

The Pillar Axiom Path Manager communicates with the Pilot over secure, encrypted XML using the TCP port 26004. If the Path Manager is installed on a SAN host, that host will require an Ethernet interface for communication with the Pillar Axiom Storage Services Manager. The network configuration must allow the SAN host to reach the Pilot management IP Ethernet interfaces.

6.2.2 Slammer Network Requirements

SAN data paths require 1 Gb/s, 2 Gb/s, 4 Gb/s, or 8 Gb/s Fibre Channel (optical) connections, which can be single or multi-mode.

The type of connection should be specified when ordering your Pillar Axiom system. Contact your Account Representative if you need to change the type of physical connection for either Gigabit or Fibre Channel.

6.3 Power-Off Requirements

If you need to turn off the system, use the Shutdown capability in the GUI to bring the system to a Shutdown state before powering off the components. Because of the redundant architecture, you might not turn off the system by switching off components (including the power distribution units).

Note: If you will be powering off the system for more than a day, remove the Slammer batteries after the system has been placed in a shutdown state so they do not discharge.

6.4 Power Cycling

Contact the Pillar World Wide Customer Support Center before power cycling a Pillar Axiom system except in the event of an emergency. In an emergency, drop all power and then contact the Support Center. Contact the Support Center before touching any power cables or switches. There are some situations where *not* power cycling the entire system is the correct action.

For *failure* testing, do not power cycle individual components without first contacting the Pillar World Wide Customer Support Center.

See also Sections 9.45, 9.50, 9.58, and 9.65.

7 Known Issues

The Pillar Axiom server issues listed in Table 7 are known at the time of this release. They are planned for resolution in upcoming releases. When available, Pillar Data Systems will provide updated software or hardware.

For additional information or help on any of the issues below, please contact your Pillar Data Systems authorized representative (see Table 3 Contact information).

Table 7 Known Pillar Axiom server issues

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
All	The UTF8 localized characters returned by the Pillar Axiom CLI and displayed by the Pillar Axiom Storage Services Manager do not always match.	Check your client setup to ensure that all proper language support is installed. This issue will be fixed in a future release.
	When the Pillar Axiom system has several hundred or more alerts, the GUI can take awhile to display the list of alerts.	Use the Pillar Axiom CLI to retrieve the list. This issue will be fixed in a future release.
	Pillar Axiom CLI allows the administrator to issue a syntactically valid "modify" command for which the parameters are the same as the existing settings in the Axiom system. In this case, Pillar Axiom CLI indicates the command completed successfully, but does not inform the user that it resulted in no actual changes to the system.	This issue will be fixed in a future release.
	In the Pillar Axiom Storage Service Manager GUI, if the administrator selects all possible objects and statistical values to display in the real time trending window, the results do not display very clearly.	Choose only the objects and statistical items that need to be displayed. This issue will be fixed in a future release.
	Cannot add recommended Storage Classes to a user-created performance profile.	This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	After uninstalling the Pillar Axiom Storage Service Manager software, log files associated with the GUI are not removed from the client workstation.	Manually delete any log files from the Temp folder in your local user path. This path depends on your OS platform. For example, in Windows XP, this path might have this pattern: `C:\Documents and Settings\<username>\Local Settings\Temp\ axiomgui<identifier>temp.zip` This issue will be fixed in a future release.
	In the GUI, the administrator must manually select each RAID controller in which logs are to be cleared.	Run a script using Pillar Axiom CLI. This issue will be fixed in a future release.

	The management software does not validate that the administrator is providing a valid netmask IP address when setting the Pilot management network settings.	Use a valid netmask for accessing the Pilot control units:
		`255.255.255.0 255.254.0.0`
		`255.255.255.255 255.252.0.0`
		`255.255.255.252 255.248.0.0`
		`255.255.255.248 255.240.0.0`
		`255.255.255.240 255.224.0.0`
		`255.255.255.224 255.192.0.0`
		`255.255.255.192 255.128.0.0`
		`255.255.255.128 255.0.0.0`
		`255.255.254.0 254.0.0.0`
		`255.255.252.0 252.0.0.0`
		`255.255.248.0 248.0.0.0`
		`255.255.240.0 240.0.0.0`
		`255.255.224.0 224.0.0.0`
		`255.255.192.0 192.0.0.0`
		`255.255.128.0 128.0.0.0`
		`255.255.0.0`
		This issue will be fixed in a future release.

	The Pillar Axiom CLI product has no option to send a request and get the response status at a later time. All requests are synchronous.	This issue will be fixed in a future release.
	The Pillar Axiom CLI product has no option to return the list of valid error codes.	This issue will be fixed in a future release.
	The event log does not report when an account has been disabled.	Check the particular account in question. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Replacing the network interface module (NIM) in each Slammer control unit (CU) within five minutes of each other can bring down the host connections.	When replacing the second NIM, verify that the data paths to the first CU are operational before beginning Guided Maintenance to replace the second NIM. This issue will be fixed in a future release.
	When a Brick RAID controller is removed during Guided Maintenance, the remaining controller may send a "RAID Controller Fault" event to the host in addition to the "RAID Controller Removed" event.	This issue will be fixed in a future release.
	On rare occasions when upgrading system software, the firmware upgrade in a Brick RAID controller can fail to complete, which causes the system upgrade to fail.	Replace the RAID controller that has failed the firmware upgrade. Then restart the system software upgrade. This issue will be fixed in a future release.
	In rare situations, when a SATA RAID controller is failing and a drive in the same Brick is taking a very long time to execute commands, a RAID controller reset may take the whole Brick offline.	Address the RAID controller failure and the drive issue separately. This issue will be fixed in a future release.
	Call Home transfers by means of an HTTPS proxy server may be reported as successful on the Pillar Axiom system when in fact the Call Home log bundle has not been received by Pillar.	Verify whether the system serial number of the Pillar Axiom system is properly registered with the Pillar Call Home server so the proxy server can login and send Call Home log bundles. This issue will be fixed in a future release.
	When an untagged Slammer port is connected to a CISCO switch that is configured to accept tagged packets, the switch drops all received packets. However, the Slammer does not report the connection problem. The Slammer simply reports that the port is Normal. Also, the port does not fail over.	Change the configuration of the CISCO switch to accept untagged packets. This issue will be fixed in a future release.
	SNMP clients receive events from the Pillar Axiom system using the specific IP address assigned to each individual active Pilot rather than the public IP address associated with the system.	The SNMP client should accept events coming from the static IP address of the active Pilot. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Pillar Axiom systems do not support the use of Windows 2003 SP1 or higher as the NTP server for the system.	For an explanation of how to provide a stable NTP for both Windows and Unix environments, refer to Microsoft TechNet article "Appendix H: Configuring Time Services for a Heterogeneous UNIX and Windows Environment" (http://technet.microsoft.com/en-us/library/bb463171.aspx). The Meinberg NTP server is a software product that can be installed on a Windows platform to provide NTP services for a Pillar Axiom system. This issue will be fixed in a future release.
	If a drive is performing marginally, the system does not notify the storage administrator until the error-rate becomes high enough that the drive is taken offline. At that point, the spare drive is used automatically.	This issue will be fixed in a future release.
	The user filename parameters of the axiomcli requests for uploading NIS files can only be filenames. If they are pathnames to the files, the requests will fail.	
SAN	The configuration of the Pillar Axiom MaxRep for SAN engine is closely coupled with the World Wide Names (WWNs) of the FC HBA ports. This means that, if the HBAs are replaced, the configuration must be redone manually.	If one must replace the FC HBA, reconfigure the replication engine. This issue will be fixed in a future release.
	There is no assurance that the labeled physical Ethernet ports on the Pillar Axiom MaxRep for SAN appliance will match the software configuration of the ports.	Log in to the appliance and check the port configuration at the OS level to get the needed information. This issue will be fixed in a future release.
	For the Pillar Axiom MaxRep for SAN appliance, the Management RMM3 port does not have an event to mark the loss of AC, a failure of the power supply, or a missing power supply.	This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	In Pillar Axiom MaxRep for SAN, RMM3 video redirection can occasionally cause the entire appliance to hang, requiring a power cycle to recover.	This issue will be fixed in a future release.
	In the Pillar Axiom Storage Services Manager, when the view of SAN host initiators is collapsed, the mapping cannot be removed.	This issue will be fixed in a future release.
	If a bad Fibre Channel host bus adapter (HBA) is installed in the Pillar Axiom MaxRep for SAN appliance, the appliance will not restart.	Replace the HBA. This issue will be fixed in a future release.
	In the Pillar Axiom Storage Services Manager, column sorting the components in the Brick hardware details view does not sort the drives in numbered order.	This issue will be fixed in a future release.
	In the SAN LUN content pane in the GUI, an administrator cannot select and delete more than one LUN or Clone LUN at a time.	Use the command line and create a script that deletes more than one LUN or Clone LUN. This issue will be fixed in a future release.
	If an iSCSI Initiator Name is configured in AIX using a string that is not in the "normalized" form defined by RFC 3722, Pillar Axiom Path Manager will not correctly associate that Initiator Name with the host entry in the Pillar Axiom Storage Services Manager.	Configure iSCSI Initiator Names in AIX using the normalized and generalized character set specified by RFC 3722 (http://www.ietf.org/rfc/rfc3722.txt). This issue will be fixed in a future release.
	Performance Profiles show a recommended Storage Class to use with a chosen profile. The GUI does not automatically switch the Storage Class field in the QoS page of a LUN when a particular Performance Profile is chosen.	Manually select a given Storage Class when creating or modifying a LUN. This issue will be fixed in a future release.
	When using a Fibre Channel HBA in a Windows environment, if a SAN error is encountered, recovery may take up to four minutes and can result in performance issues.	Inspect your SAN configuration for bad connections and resolve any issues that may exist. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	A defective Fibre Channel network interface module (NIM) may result in a kernel panic and a warmstart of the Slammer control unit.	Replace the defective NIM. This issue will be fixed in a future release.
	On a Linux host, paths may not be restored to an online status after a warmstart or path failure.	Run the Qlogic or Emulex Rescan utility. This issue will be fixed in a future release.
	Using a software iSCSI initiator when an issue in an ESX server configuration (misconfigured LUN) exists can cause performance issues.	Be sure the ESX server is correctly configured or switch to a hardware iSCSI initiator. This issue will be resolved in a future release.
iSCSI	In rare circumstances, iSCSI systems under heavy load or with high numbers of network errors may warmstart.	Try to rebalance the load. If you're getting a large number of network errors, replace the iSCSI network interface module in the Slammer. This issue will be resolved in a future release.
	On rare occasions, the iSCSI network interface card can hang, which causes the system to warm start.	Replace the iSCSI card, or contact Pillar World Wide Customer Support for assistance. This issue will be fixed in a future release.
	During Windows startup, the iSCSI Software Initiator may attempt to register with the Microsoft iSNS Server v3.0 if it has been configured to do so. If the iSNS Server is installed on the same host and has not started yet, the iSCSI Initiator may fail to register with the server.	Edit the Windows registry to add a dependency that causes the iSCSI Initiator to wait for the iSNS Server to start. This issue will be fixed in a future release.

8 Resolved Issues

A number of issues, some previously undocumented, have been resolved in this release. Items that were documented as known issues in the previous release and are resolved with this release are described below. These items are no longer product issues.

Tip: For a complete listing of resolved issues, please contact the Pillar World Wide Customer Support Center (see Table 3 Contact information).

Table 8 Resolved Pillar Axiom server issues

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
5.2	All	In the Pillar Axiom Storage Services Manager, the Tools menu contains no items, even though it appears to be active.
		The system keeps track of historical statistics. Even though no I/O activity over a given set of Slammer ports is currently occurring, statistics will exist from a historical standpoint.
		Sometimes, the Pillar Axiom Storage Services Manager GUI consumes too much memory and CPU cycles after being left running for several days.
		When running multiple Pillar Axiom Storage Services Manager sessions that access different systems, if the administrator logs off from one system, the other session shows the last system to be logged off in the dialog window when the user attempts to logoff from a different system. Even though the dialog text is incorrect, the administrator is logged off the other system.
		Resetting the Support Administrator credentials by means of the USB key reset procedure is not supported.
		The Pilot may fail to restart following a power outage if the Pilot was writing to the drive at the time the power failure occurred. This is caused by a single, uncorrectable sector on the Pilot drive and the following message appears: "fsck.ext3: Input/output error while recovering ext3 journal of /var."
		Replacing the network interface module (NIM) in each Slammer control unit (CU) within five minutes of each other can bring down the host connections.
		The Pilot uses a version of OpenSSH (OpenBSD Secure Shell) that has several known vulnerabilities.
	SAN	During LUN creation, the Storage Class field does not default to a specific Storage Class in the Quality of Service (QoS) tab.
		When upgrading from release 4.x to release 5.0, even though the Quality of Service (QoS) settings for the LUNs are correct, all LUNs show "Custom" as the performance profile.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		When an administrator moves a SAN host from one host group to another host group, the host may not get associated with the new host group.
		When a SAN host is moved into a host group, sometimes a remapping occurs and data access is lost when the new mappings are applied.
		In the Pillar Axiom Storage Services Manager, when creating a host mapping for a LUN in the mapping dialog tab for the first time, the newly mapped host does not show up in the list of mapped hosts.
		When there are global mappings, the GUI displays mappings for a given SAN host individually and under a Host Group.
		The GUI displays multiple error dialogs referencing previously failed scheduling creation attempts.
		In the Pillar Axiom Storage Services Manager, double-clicking a report that is listed in the Generated Reports content pane brings up the "Generate Report" dialog. Double-clicking should bring up the download window.
		The help button for the Trending window in the Pillar Axiom Storage Services Manager does not work.
5.0	All	Due to an incompatibility issue between the firmware on legacy Fibre Channel (FC) Bricks in releases below 4.3.0 and release 5.0, all system upgrades to release 5.0 need to first be upgraded to release 4.3.15 or higher as an intermediate step. If this intermediate step is not performed, any legacy FC Bricks on the system end up Critical with the possibility for data loss.
		In certain circumstances, a Fibre Channel Brick can get a stuck command on one RAID controller when a Pre-Emptive Copy (sometimes called *copy-away*) completes. This situation can result in I/O to that controller being halted, which can result in logical volumes going offline.
		System Error page displayed after staging software update package that contains a drive firmware update.
		In releases prior to 5.0, new factory Bricks added to a system would be accepted without requiring confirmation of the "add". As of release 5.0, when one attempts to add factory Bricks, what happens with regard to prompting depends on whether the Storage Domain feature is licensed. If it is, the admin is prompted for the Storage Domain into which the Brick should be added. This happens even when only a single Storage Domain is configured, in case the admin wants to create a new Storage Domain and add the new Brick to that domain. If the Storage Domain feature is not licensed, the new Brick being added is automatically accepted without requiring confirmation.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		Multiple concurrent restore operations on the same volume from a clone of that volume do not work. The first restore may succeed, but the second hangs and be outstanding forever. The situation gets more complicated if a Slammer control unit fails while the restore operations are in progress.
		There is no event in the event log indicating that shutdown had succeeded.
		Under a specific heavy workload pattern, a Brick RAID controller can sometimes run out of resources, resulting in I/O to that controller being stuck for a certain amount of time. This can result in pinned data and the associated LUNs alternating between Online and Conservative.
		Resetting the Support Administrator credentials by means of the USB key reset procedure is not supported.
		During very short (<200 msec) AC power disruptions, a Slammer control unit (CU) can sometimes go offline and stay offline.
		When a Brick has a Power supply replaced, it is possible for the serial number of the newly replaced power supply to be incorrectly reported in the GUI.
		When a Brick experiences a read error on two drives in the same stripe or a read error during a RAID array rebuild, the data for the associated stripe is lost and a Bad Block Log entry is created to track the bad stripe. If any reads are sent to this stripe, they are rejected with an error. If the stripe is written, the Bad Block Log entry will be cleared.
		The Pilot uses a version of OpenSSH (OpenBSD Secure Shell) that has several known vulnerabilities.
	SAN	When mapping a LUN to Host Groups, the GUI lets the user attempt to change the LUN number mapping for a given host in the Host Group, although this is not allowed. All hosts in a given Host Group receive the same LUN number mapping.
		A Pillar Axiom system might experience a software fault if configured to enable iSNS, but with an invalid iSNS server IP address.
		If an iSNS Server has been registered on a Pillar Axiom system and the iSNS Server cannot be contacted by either static IP address or DHCP, the system may not be able to fully shut down. If a shutdown operation is requested, the task completion operation may stop at less than 100%. This can prevent the system from performing major release software upgrades.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		In releases prior to R5, one could define a separate Slammer port masking configuration for each LUN-host combination. In R5, Slammer port masking becomes a property of a SAN LUN, and the same port mask will be in effect for all hosts accessing this LUN. When upgrading from R4 to R5, the single Slammer port mask that will be created for each given LUN is a logical "OR" of all masks in effect for that LUN in R4. So, for a given LUN, if a Slammer port was enabled for any host in R4, that port will be enabled for all hosts after the upgrade to R5.
		The Brocade HBA may not function properly when connected directly to the Pillar Axiom system. The HBA appears to be connected; however, no LUNs are visible on the host.
4.3	All	During a warmstart, attempts to reinitialize the internal Ethernet controller can sometimes fail, which results in a Slammer control unit (CU) failover.
		Unable to parse and see Brick S.M.A.R.T. data using the Pillar Axiom Statistics Tools.
		In rare cases, the core dump may be unavailable after the Slammer warmstarts.
		Relevant Brick logs might not be automatically collected for Call Home when a fault event for a RAID controller occurs.
	SAN	When isolating a replication pair using the command `axmrepsan isolate_replication_pair` ***`pair_name`***, sometimes the replication pair can go into the AWRY state, which is verified in the output generated by the command `axmrepsan get_replication_pair_status`.
		If there were a large number of SAN cache de-stage requests pending, an incoming write from a host could be delayed long enough that the host would abort.
		Due to a software race condition, SAN LUNs could fail to come back online after a system power failure.
4.2	All	Following the addition of Bricks and code upgrades, sometimes the system inadvertently turns off Brick event monitoring. As a result, the system is not aware of problems and of error recovery actions being taken by the Bricks.
		When a Slammer control unit failback fails and a second failback succeeds, the battery-backed write cache may not be properly established for logical volumes, causing a significant performance loss.
		Due to an internal locking issue the Pillar Axiom system may warmstart in overloaded conditions.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		When an external security audit is run against the Pilot management controller, sometimes the audit indicates a security violation (the server supports the TRACE and/or TRACK methods).
		A failing drive on a Fibre Channel Brick may cause the Brick to become unresponsive to I/O requests for a few minutes if a write command fails to the drive. This lack of response may cause LUNs and filesystems to go offline temporarily.
		When a Pillar Axiom system experiences several Slammer control unit failovers in a short period of time, sometimes the system stays in write-through mode after failback, causing performance to degrade severely.
		A Slammer control unit may warmstart unexpectedly after performing management operations (such as creating and deleting Clone LUNs) for several hours.
		After a successful shutdown, sometimes a restart can result in the system erroneously reporting lost blocks.
		Pillar Axiom 300 systems may not be able to complete a warmstart successfully, leading to a Slammer control unit failover or failback when the warmstart times out.
		Discovery of a triply redundant volume does not produce a configuration-on-disk (COD) error as expected.
		When a Pilot control unit (CU) is replaced in a Pillar Axiom system, sometimes the active Pilot CU fails to update some of the components on the replacement Pilot CU correctly.
		Events are generated for non-critical internal I2C bus errors.
		Not able to download the Pillar Statistics Tools for Windows and Linux platforms.
		During a Slammer failover operation, an additional Slammer warmstart may result.
		System failure may result from a user application attempting to write to a volume that is in the process of being thinly provisioned.
		If a Slammer control unit (CU) warmstarts while the system configuration is being changed, the CU may subsequently warmstart again.
		If one creates a thinly-provisioned volume, and then makes it fully provisioned by changing the current capacity to equal the maximum capacity, sometimes an error can occur (because of internal rounding errors).

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		A memory leak in the message buffer may cause a Slammer control unit to warmstart unexpectedly during routine management operations that last several hours.
		When deleting an extremely large LUN or filesystem, the Brick takes a large amount of time to complete the transition of the related storage from allocated to unallocated. During this time, the Brick is temporarily inaccessible until the operation completes.
		During failover and failback, a Slammer control unit may warmstart unexpectedly.
		A Slammer control unit (CU) may fail when intense I/O on a sparse volume causes a high rate of infill allocations in a short time. This high rate of allocations can cause a count mismatch, resulting in CU failure.
		During a software update, you may receive an alert indicating that the Update Software Task Failed. This failure may be the result of some internal events that may have not been forwarded (or processed) during a Slammer control unit warmstart.
		PIM replacement, which includes powering off and powering on of the parent Slammer control unit (CU), may cause NonOptimizedAccess events for APM.
		The `GetStorageConfigDetails` PDSCLI command returns raw capacities based on a RAID 5 geometry. This does not account for capacity requirements when using Distributed RAID.
	SAN	During replication, a background data migration deadlock can occur that causes a Pillar Axiom system to warmstart.
		On rare occasions, the Pillar Axiom system has problems tracking SAN host logouts and logins, which can result in a warmstart of a Slammer control unit.
		Migrating a currently replicating LUN to another Storage Class may result in a system warmstart.
		Unable to save a change only to the access bias of a LUN. Additional changes to the LUN were required for the save to be successful.
		The SAN Remote Replication feature requires both a Remote Replication license and a Clone LUN license to function. The fix requires only the Remote Replication license.
		Unexpected errors may occur when attempting to move LUNs between volume groups. The system may think that there is insufficient capacity in the volume group for which to move the LUN. Also, the reporting of these errors may state that it is an "Xpath error" which can be ignored.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		During the checkpoint process, a SAN Replication pair may go AWRY.
		When SAN Replication pairs are running in two directions at the same time (source to destination and destination back to the same source), the pairs may stop replicating, eventually causing the pairs to go AWRY or ISOLATED.
		When SAN Replication pairs are running in two directions at the same time (source to destination and destination back to the same source), the pairs may stop replicating, eventually causing the pairs to go AWRY or ISOLATED.
		During replication, the Pillar Axiom system might experience multiple software faults, leading to a full system restart.
		When you create a SAN Replication checkpoint for which the source is ready but the target is not, numerous alerts stating that the target checkpoint failed are issued.
		When the primary and secondary Pillar Axiom systems hosting a replication pair are started at the same time, sometimes one or both systems may not complete the restart, which can result in Pilot failover and a significant delay in getting both systems operational.
		During SAN replication operations, if failover-failback occurs, additional warmstarts might occur, disabling the Slammer control unit.
		In a clustered environment that uses SCSI persistent reservations, a SAN host that has been rebooted and lost its persistent reservation or registration for a cluster LUN (which had LUN mapping enabled) can still access the LUN.
		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.
		When AxiomONE Replication for SAN replication pairs are replicating, numerous Event Out Of Sequence events can appear in the event log.
		When the owning Slammer of a Clone LUN is in a non-normal status, the health of the Clone LUN may incorrectly be displayed as Online when in fact it is Conservative.

9 Additional Notes

For items in this section that refer to inserting and/or removing field replaceable units (FRUs), please refer to the *Pillar Axiom Service Guide* for more information.

For items in this section that refer to provisioning or configuring a Pillar Axiom system, please refer to the *Pillar Axiom Administrator's Guide* for more information.

9.1 Host Queue Depth on SAN Hosts

The recommended maximum 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 the Pillar Axiom system. Exceeding this value may cause I/O errors if the input/output queue of the Pillar Axiom system is exceeded.

This value is typically set in the BIOS or similar firmware configuration of the HBA on the SAN host. Consult your HBA documentation for the setting that controls the maximum I/O queue depth for your HBA and for configuring this setting.

9.2 Limitation of the Linux sginfo Utility

The `sginfo -l` utility (part of the Linux sg3_utils package) has a limitation by which it can only display up to 31 LUNs. To display the actual number of devices recognized by a host, use the fdisk -l utility instead.

9.3 LUN Ranges in Windows

Windows 2000 and 2003 will not configure LUN 255. If you configure a LUN in the Slammer at address 255, Windows will not see the LUN.

9.4 Issues with LUN Capacity Calculations on Solaris

The Solaris operating system calculates the size of a LUN using disk geometry information from Mode Sense queries rather than the more common and accurate practice of using the response to a Read Capacity Query. For Pillar Axiom LUNs larger than approximately 400 Gigabytes, this calculation can result in a reported capacity that is different from the Pillar Axiom configured value.

The Solaris format utility may return an error stating that it is adjusting the number of sectors on the Pillar Axiom LUN or may indicate that the number of heads is something other than 64 or that the number of sectors is something other than 128 when Solaris adjusts the number of cylinders to be 65,533 during the size calculation. If format returns an error, it is typically:

> Mode sense page(3) reports nsect value as 128, adjusting it to 127

Disk geometry information does not apply to SAN LUN arrays on Pillar Axiom systems. This information is returned, however, in Mode Sense with the number of heads and sectors being 64 and 128 and with the number of cylinders varying for those operating systems (such as Solaris) that calculate LUN size rather than using the actual Capacity.

If the difference between the information calculated by Solaris and the actual LUN size is an issue for your applications, create and use a unique disk label or /etc/format.dat entries for the Pillar Axiom LUNs.

9.5 VMware ESX Server 3.0.1 Connections to Pillar Axiom iSCSI Systems

When booting from SAN, only one path to the Pillar Axiom system should be configured in the iSCSI HBA BIOS. The boot LUN is assigned a LUN ID of 0 (zero) to which the iSCSI adapter ports must be mapped.

9.6 Avoid GM of Slammer Components During Heavy I/O Loads

During the Guided Maintenance (GM) of Slammer CU components, the system attempts to flush all pending I/O writes to storage so that the system can go into write-through mode. Going into write-through mode ensures that no data is in cache should a catastrophic error occur. If there is heavy write activity during this time, an alert may be generated in the GUI warning the administrator that the system cannot properly prepare for the upgrade.

In this case, retry the operation or queisce the hosts that are writing data to the Pillar Axiom system.

9.7 MS iSNS Server Could Add or Remove Discovery Domain Members Quietly

Under certain circumstances, the Microsoft iSNS Server v3.0 may add or remove members of the Pillar Axiom's discovery domain without notifying the Pillar Axiom system. If iSNS access control is enabled, the missing notifications can cause the Pillar Axiom iSCSI target to accept or reject iSCSI initiator logins when it should not.

To prevent this problem from occurring, follow these guidelines:

- When creating a new discovery domain, add the Pillar Axiom system to the discovery domain before adding any iSCSI initiators.
- Disable a discovery domain set before deleting it.
- Ensure that an iSCSI initiator is registered with the iSNS server before adding it to an existing discovery domain.

If a problem already exists, any of the following actions will cause the Pillar Axiom system to query the iSNS server for the latest discovery domain information:

1. Disable and re-enable iSNS server registration in the Pillar Axiom system.
2. Disable and re-enable the discovery domain set(s) in the iSNS Server GUI.

9.8 HP-UX HBA Connections to Pillar Axiom Systems

The Pillar Axiom user interfaces show that host Fibre Channel (FC) HBA ports are either Connected or Not Connected to the Slammer ports. The meaning of Connected is that the HBA port on the SAN host has logged in to the port on the Slammer using the FC protocol.

In most operating systems, host ports log in to the Slammer ports immediately after the two are physically connected and enabled and remain logged in until the physical connection is broken. Therefore, Connected in the UI effectively means that there is an enabled physical connection between the ports.

Some HBA device drivers on HP-UX, however, use a different approach—they log out from the connection when there is no traffic to send. An HP-UX HBA port often shows as Not Connected even though there is an enabled physical connection between the ports.

If a SAN host has HP-UX initiator ports and HP HBAs, when associating that host with a LUN, use the HP-UX Compatibility Mode option. When this option is enabled, the system determines LUN numbers using the HP-UX addressing scheme, allowing up to 255 LUNs. Also, when enabled, the host cannot have a visible LUN using ID 0. You can verify the current host mappings in the Pillar Axiom Path Manager tab.

9.9 LUN Assignment and Accessibility

If you use Pillar Axiom LUN masking or switch zoning and do not use LUN assignment, you may create a situation in which a LUN is not exposed on the ports on which you want to access it. To avoid this situation, it is recommended that you assign the LUN to the Slammer control unit (CU) on which you have the mapping set.

9.10 LUN Assignment When the System Restarts

When an administrator creates a LUN and chooses auto-assign, the system decides the optimal control unit (CU) on which to place a LUN. The LUN resides on that CU from that point forward, even after a system restart, unless a non-optimized access (NOA) event occurs or the administrator reassigns the LUN to a different CU. If the administrator later modifies the LUN and assigns it to a different Slammer CU, the LUN loses its auto-assign status.

Note: As of release 5.0, the system no longer re-balances LUN assignment during a system restart operation.

9.11 Blacklisting Local Drives on RHEL4 Platforms

Device Mapper on RHEL4 U4 platforms may display local SCSI SAS or SATA drives along with the FC drives as multipathed. Including local drives can be avoided by blacklisting the devices in the `/etc/multipath.conf` file:

```
devnode_blacklist {
      wwid 26353900f02796769
      devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st|sda)[0-9]*"
      devnode "^hd[a-z][0-9]*"
      devnode "^cciss!c[0-9]d[0-9]*[p[0-9]*]"}
```

The above work-around is suggested by Redhat in their knowledgebase at http://kbase.redhat.com/faq/FAQ_85_7319.shtm.

After running the following commands, the local disks should no longer be listed in the new multipath maps:

```
multipath -F
multipath -v2
```

9.12 iSCSI Software Initiator May Have Two Names Associated With It

The Microsoft iSCSI Software Initiator may sometimes use an iSCSI Initiator Name other than the one set in its configuration. For example, if the configured Initiator Name ends with the Fully Qualified Domain Name of the host, when making iSCSI connections, the Software Initiator may use a Name ending with only the node name of the host. In this case, the Pillar Axiom Storage Services Manager GUI and CLI will report that the host is using two iSCSI Initiator Names, both the configured name and the name it is actually using.

9.13 Resetting the Primary System Administrator Password

If you forget the Primary System Administrator password, you can reset it in these ways:

- Use a Type 1 Administrator account, if one exists, to reset the password. A Support Administrator cannot reset the Primary Administrator password.
- Contact the Pillar World Wide Customer Support Center for the encrypted file (for resetting the password), which may be placed in a USB key. Use the USB key as instructed.

It is strongly recommended that you set up an additional Type 1 Administrator account when you install the system. A Type 1 Administrator can modify account passwords without knowing the previous password for any accounts.

9.14 Uploading Software Update Packages over Slow Connections

It is not recommended that you upload a software update to your Pillar Axiom system over a slow connection (such as a WAN connection). Use an internal network connection (10 Mb/s or greater) only.

9.15 When Updating the Software

Whenever you update Pillar Axiom software, ensure that all non-Pillar Data Systems components are working correctly with all redundant paths enabled and no maintenance being performed on any other component in the network.

9.16 Software Update Process Waits for all Tasks to Complete

If a system software update is initiated while Pilot tasks are in progress (such as an automatic Call Home and log collection), the update process will wait, and no information is provided stating that the update is waiting for the tasks to complete.

Here are some suggestions to prevent the software update process from going into a wait state:

- Schedule updates when it is known that long-running tasks will not be running.
- Before starting a software update, be sure all tasks are complete.
- During the software update, review the tasks that are running. However, before you cancel those tasks, contact the Pillar World Wide Customer Support Center for assistance in determining which, if any, tasks can be canceled.

9.17 Non-Disruptive Software Updates

The Pillar Axiom system implements non-disruptive software updates by warmstarting the Slammer control units (CUs) and restarting the Pilot CUs to bring up the new software. As each Slammer CU warmstarts, there is a temporary protocol service disruption of a few seconds on each CU. This disruption is typically non-disruptive to most applications and protocols.

For SAN Slammers, if the HBA timeouts and retries are set correctly, this brief protocol disruption should be handled gracefully by most operating systems and applications.

However, any application or operating system that bypasses the Fibre Channel protocol stack and issues SCSI commands with short timeouts may not be capable of handling the brief interruption of a non-disruptive software update.

9.18 WWN Designation Changed in Release 2.0

Starting with the 2.0 release of Pillar Axiom 500 systems, the WWN was changed to use a common base World Wide Node Name (WWNN):

- In release 1.x, each Slammer was assigned a unique WWNN with the Slammer Fibre Channel ports being assigned World Wide Port Names based on the WWNN of the Slammer. For each Slammer, CU0 would have World Wide Port Names using 1 and 3 and Slammer CU1 would have World Wide Port Names using 2 and 4 to indicate the port, based off the Slammer WWNN.
- Starting with release 2.0, the entire Pillar Axiom system has a single base WWNN based on the MAC address of Slammer CU0. The World Wide Port Names are derived by a fixed formula from this single base WWNN using the Slammer, Slammer CU, Slammer Port Number, and Port Type.

Starting with release 4.3, the Fibre Channel (FC) World Wide Name (WWN) for the Series 3 Slammer that contains a SAN 8 Gb/s FC network interface module (NIM) is derived from the 8 Gb/s FC HBA itself. The WWNs for both Slammer CU ports are printed on a label located on the faceplate of the HBA. They are also displayed in the GUI.

Important! If you replace one of these NIMs, you might need to rezone your FC switch or change the configuration of any SAN replication devices to account for the new WWNs of the replacement NIM.

9.19 Linux 2.6 Marks Filesystems Read-Only When Access to Boot LUN Is Lost

Linux 2.6 is sensitive to losing all paths to a boot LUN. When a SAN host loses access to all paths to a boot LUN, to prevent data corruption, the Linux system marks the file system on the client as read-only. This behavior is as designed and would occur regardless whether the Pillar Axiom Path Manager is installed on the SAN host. To improve the path recovery time, Pillar recommends that you:

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

9.20 Dynamic LUN Expansion Not Well Supported on Linux

In general, device-mapper does not support dynamic LUN expansion. Specifically, the latest QLogic rescan utility on a Linux host does not gracefully handle LUN expansion on a Pillar Axiom system.

If you expand a LUN on the Pillar Axiom system, you need to reboot the Linux host to make the LUN expansion visible.

9.21 Status of Logical Volumes

System health screens in the GUI display the status of hardware and firmware components of the Pillar Axiom system. The overall system status icon on the bottom of the screen is a summary of the hardware status and does not reflect the status of the logical volumes.

A hardware problem typically causes logical volumes to go offline or to a degraded state. Because this is not always the case, you should check the state of the volumes or any associated alerts that may be listed.

9.22 Changing the Time on a Pillar Axiom System

Even though your environment may not require an NTP server, Pillar recommends that you use an NTP server.

Note: If an NTP server is controlling the time on a Pillar Axiom system, you should change the date and time *only* on that time server.

If you are not using an NTP server, we recommend not changing the date once the initial installation is complete and the system is operational. If you change the date on a Pilot or Slammer by more than 15 minutes, the NTP daemon will mistrust the request and exit. Changing the date may result in reporting of events and alerts with bad dates. Should you do this or see date stamps on events or alerts that are obviously invalid, contact Pillar the Pillar World Wide Customer Support Center for recovery assistance.

Tip: If you are about to switch to using an NTP server, be sure the current time on the Pillar Axiom system is within 15 minutes of that on the time server; otherwise, a Pilot failover may result.

Tip: When initially setting up NTP and TimeZone on your Pillar Axiom system after installation, restart the system to make sure NTP on all internal components is properly synchronized.

9.23 Keeping Clone LUNs From Being Deleted

When the repository for Clone LUNs (formerly called Snap LUNs) consumes more than 90% of its allocated capacity, the system is likely to begin automatic deletion of Clone LUNs. When repository usage crosses this 90% threshold, the system creates a system alert SnapLUNStorageFillingButCannotGrow ("Extra space for Clone LUNs has reached maximum and is nearly full") to warn of this possibility. If you see this Administrative Action, you should manually delete some of the Clone LUNs.

If you want to use lots of I/O on a Clone LUN, to keep the Clone LUN from being deleted, you should allocate a size of 120% of the source LUN. For multiple Clone LUN descendents of a source LUN, each one requires more space, so you should allocate an additional 50% for each Clone LUN that you intend to have in existence at a given time. Actual storage space used for the repository is only grown to the amount of space being used.

Clone LUNs are not intended for heavy I/O, they are intended to be temporary—anywhere from minutes to several weeks in existence. As long as they are deleted the space will be recycled. All repository space is recycled when the last Clone LUN is deleted.

9.24 Deleting Clones From the Youngest to the Oldest

When several clones are deleted at the same time using the GUI, the process can take about seven minutes for each clone. The length of time it takes to delete a clone is related to the time it was created and how it relates to the clone storage space being used. The best approach to deleting clones would be to choose the youngest clone first and work your way back to deleting the oldest.

Deleting a clone of the most recent data may take a long time because this action requires reassignment of data to ensure that the downstream clones (if any exist) will have consistent data.

Deleting a clone of the oldest data is done immediately because there is no such reassignment necessary.

9.25 Small Maximum Clone LUN Space

When creating a LUN, you can specify Max space for Clone LUNs that is as little as 50% of the LUN capacity, or lower (minimum 1 GB). Choosing a small number is only appropriate for LUNs that have the following characteristics:

- It has only a few Clone LUNs at any given time.
- Its Clone LUNs will have short lifetimes (for example, they will be deleted after making a backup).
- It will get minimal write activity while there are Clone LUNs.

Specifying a larger Max space for Clone LUNs (for example, up to 300% of the LUN capacity) is safe and reasonable. The system will allocate a small fraction of the specified Max and will increase the space automatically as warranted by Clone LUN activity up to that Max.

It is good practice to delete Clone LUNs when you are done with them. Old Clone LUNs run the risk of running out of space and losing synchronization with their source LUN. If that occurs, their data will be corrupt, and they will be automatically deleted. If you need a long-term copy of an active LUN, consider using the Backup to Disk option instead.

9.26 Reassigning a Logical Volume to Another Slammer or Control Unit

If you reconfigure the Slammer or the Slammer control unit (CU) to which a logical volume is assigned, the Pillar Axiom system stops the resource and moves the volume to the new location. Attempting to use a logical volume while it is being moved can result in lost data.

Note: You might find it prudent to stop host I/O before moving a LUN if a manual move is slower than an automatic non-optimized access (NOA) move.

Important! Pillar recommends that all SAN clients unmount the LUN that is to be reassigned before reassigning the logical volume.

9.27 System Scaling

You may increase storage capacity on a system by adding components and increasing assigned capacities. You may not decrease storage capacity of a system without the help of a Pillar Data Systems authorized representative.

9.28 Information Screens for Slammer Power Supplies

In the System Health screens for the Slammer Components, the information fields for Slammer power supplies are intentionally blank.

9.29 Cause of Degraded Status for a Volume Has Changed

Beginning with release 4.0, a failed, missing, or pulled drive will not by itself cause a logical volume to become degraded. Only an inability to write to the mirror in a doubly redundant volume will result in a degraded status.

The inability to write to the mirror can occur because a Brick that underlies that mirror is not available or is not communicating, or a RAID LUN has faulted.

9.30 A Drive May Display Blank Data in the GUI or Create an Alert

Drives are validated by the Pillar Axiom system. In some cases, the following may occur:

- The GUI may report a blank part number or serial number for a drive and, occasionally, a "Cannot read" status for that drive. However, the system will perform normally.
- The system generates an Incompatible Hardware alert. In this case, contact the Pillar World Wide Customer Support Center for assistance in replacing the drive or resolving the alert.

9.31 Replacing Hardware Components

Use the Guided Maintenance feature of the GUI when replacing hardware components. Guided Maintenance provides instructions to you and performs tasks to get the system ready for the component replacement. Be sure to follow the instructions provided.

Notes:

- Adding memory to existing Slammer CUs in the field is not supported in this release. Contact your Account Representative for assistance.
- Replacing Slammer chassis are not supported in this release.

  If you need to replace a Brick chassis, contact the Pillar World Wide Customer Support Center for assistance.

9.32 Alternative to Guided Maintenance Identify Step for Pilots

Do not rely solely on the Pilot Identify process during Guided Maintenance. That process uses the hard drive LED as the method to identify the selected pilot and, depending on the activity on the Pilot control unit (CU), it may not be possible to identify clearly which CU is being beaconed.

The best method to identify a Pilot CU is to match the serial numbers reported in the GUI with the labels on the Pilot CUs.

9.33 Pilot CUs Must Not Be Powered On Independent of Replacement Procedures

After receiving a replacement 1U Pilot control unit (CU), do not power it on outside of the Pilot replacement procedure documented in the *Pillar Axiom Service Guide*. If a Pilot CU is powered on prematurely, you must contact the Pillar World Wide Customer Support Center. Also, when you need to replace a Pilot CU, contact the Support Center for assistance.

9.34 Reuse of Pilot Control Units Can Cause Problems

When removing a Pilot control unit (CU) from a Pillar Axiom system, mark that Pilot CU so it can be clearly identified. Once a Pilot CU has been removed from the system, never put that Pilot CU back into the same Pillar Axiom system without assistance from the Pillar World Wide Customer Support Center.

CAUTION: Do not move a Pilot CU from one Pillar Axiom system to another; otherwise, data loss may result.

Important! Always use only Pilot CUs that have been shipped from Pillar logistics or manufacturing.

9.35 Differentiate Between FC RAID Bricks and FC Expansion Bricks.

The system must be able to tell one Fibre Channel (FC) Brick from another. The thumbwheel in the ES component at the back of a FC Brick enables you to make this distinction. As such, you must set the FC RAID Brick thumbwheel to 0 and the FC Expansion Brick thumbwheel to 1.

9.36 Replacing a Drive

When replacing a drive, always use a new one from Pillar Data Systems.

- When replacing a drive, wait at least 30 seconds after removing the old drive before inserting the new drive.
- Do not reseat a drive unless instructed to do so by the Pillar World Wide Customer Support Center.
- Do not attempt to replace a failed drive with one from another Brick or from another Pillar Axiom system.
- If testing Drive Pull, wait a few seconds after removing the drive before reinserting it. Be sure to check for alerts to accept the drive.

  Important! You should contact the Pillar World Wide Customer Support Center before pulling a drive for test purposes.
- If a drive fails to be accepted into a Brick and the drive is set to Rejected status, do not attempt to use that drive. Contact Pillar Data Systems for another drive and for assistance.
- If an alert asking you to accept the drive is generated, be sure to select the Accept Drive option, which in some cases will initiate a copyback operation.

Important! If an alert to Accept a Drive is ever answered negatively, do not attempt to use that drive again. Contact Pillar Data Systems for another drive.

Contact the Pillar World Wide Customer Support Center for a new replacement drive.

9.37 Moving Drives

Do not move drives from their original positions. If you move a drive, all data on that drive will be lost. If multiple drives are moved, you will lose data.

If a drive is defective, use Guided Maintenance in the Pillar Axiom Storage Services Manager GUI to replace the drive.

9.38 Reseat Drives before Powering On New or Replacement Bricks

The drive latch may appear to be fully latched, but sometimes the drive is not making good contact with the Brick chassis midplane. With poor contact, the drive will fault and its status will be displayed as Critical (because it is missing).

To prevent loose drives and as a precaution, before powering on a new or a replacement Brick, visually inspect each drive to verify that they are fully seated.

If a drive is not fully seated, either or both of the following will be true:

- The metal portion of the carrier will be visible.
- The front of the drive carrier will not be flush with the other carriers.

To seat an improperly seated drive, perform the following steps:

1. Press on the drive to latch them.
2. Press the drive carrier firmly until it snaps into place.
3. Snap shut the latch to lock the carrier in place.

Important! Do not unlatch and re-latch a drive carrier unnecessarily. Doing so can lead to potential troubles in the future.

9.39 Testing Multiple Drive Failures

Important! Do not test multiple drive failure scenarios in the same Brick storage enclosure without contacting the Pillar World Wide Customer Support Center for guidance.

9.40 ACT LEDs on Drives Can Blink When Inactive

When there is no I/O activity on a Brick storage enclosure, the RAID firmware runs a background operation that scans all drives for media errors and, if media errors are found, performs repair operations. This background activity causes the ACT LEDs to blink green on the idle system or Brick. Such activity can take several hours to complete. When host I/O resumes, this background operation stops; it resumes only when there are no further I/Os from a host.

9.41 Replacement of Brick Storage Enclosures

To avoid data loss, contact the Pillar World Wide Customer Support Center before you attempt to replace an entire Brick storage enclosure or Slammer storage controller. The Support Center can help you determine whether a particular logical volume is physically on the Brick.

9.42 Adding a Brick Generates Error and Warning Messages

When you add a Brick storage enclosure to a Pillar Axiom system, the system begins the process to bring the Brick online. While the system is bringing the Brick online, you will see the message: Topology Discovery Task.

Also, you may see a series of error and warning messages similar to these:

- Fibre Channel RAID Array Inaccessible
- Fibre Channel Path to Brick Failed
- Software Update Succeeded

These messages are normal and to be expected.

The system creates a system alert when the Brick is powered on. This alert allows you to assign the Brick to a particular Storage Domain.

During the bring-up process, the status of the Brick goes from red to yellow to green. After the system completes the process, the Brick shows a Normal status and removes all alerts related to adding the Brick. If any alerts remain, contact the Pillar World Wide Customer Support Center.

9.43 Testing RAID Rebuild and Simulated Drive Fault

You can use Guided Maintenance to identify a Brick and to show the location of an individual drive for testing drive pulls. But the “Prepare System” and “Replace Hardware” functions should not be used when testing or demonstrating RAID rebuild and drive replacement where the existing drive is to be removed and then re-installed.

The Guided Maintenance process is intended for use only when a drive, or other FRU, has encountered a fault and is to be replaced with a new drive or other FRU. If Guided Maintenance “Prepare System” and “Replace Hardware” is used to replace a drive, you will be instructed to remove and replace the drive. The Pillar Axiom system may defer any further actions on the Brick until this is done, which may result in the Brick Redundancy repair actions not being initiated properly.

To test Drive Fault, simply pull the drive. Wait a few seconds until drive activity is observed on either the top or bottom drive LEDs on all other members of that RAID array, then carefully reinsert the drive and make sure it is fully seated and latched in place. The background tasks to rebuild the array and then copyback the array data to the re-inserted drive should start automatically and be displayed within a few minutes, depending on overall system activity. If an alert to accept the drive is displayed, be sure to select “yes” to accept the drive.

If a drive is genuinely faulted, use the Guided Maintenance menus to identify the Brick. Note the position of the drive in the Brick, prepare the system for replacement, and then click **Replace Hardware** to remove the old drive and replace it from spares as instructed.

9.44 Brick Issues Can Cause a Slammer to Warmstart

The Array Manager software component in Pillar Axiom Slammers requires a quorum of successful I/O for its metadata processing. Sometimes a quorum cannot be met because some Bricks return a Busy state due to underlying issues on the Bricks. After a number of unsuccessful retries, the Array Manager can fail a health check and cause the Slammer to warmstart.

In these cases, when you resolve the underlying Brick issues, the array manager will be able to successfully access metadata. To help avoid Slammer warmstarts, ensure that all drives are functioning and Bricks remain operational.

9.45 Testing Brick Power Loss

Pillar Professional Services can assist in testing power loss to Brick storage enclosures.

9.46 Use Care When Recovering a Faulted Data LUN in a Brick

In any given Brick storage enclosure, if there are enough drive failures that cause a data LUN on that storage enclosure to fault, care should be used in recovery.

When the drives are replaced, as soon as enough drives become available to allow the Pillar Axiom system to perform a `RecreateRAIDArray` task, the system generates an alert.

CAUTION! As a safety measure, contact the Pillar World Wide Customer Support Center before accepting the alert. Because the Brick being recovered may contain persistence data (system configuration), accepting the alert might cause system configuration information to be lost, which could cause data loss even for resources not physically on the Brick in question. (See also the **Important** notice in Section 9.48.)

If you do not accept the alert, in most instances the Support Center should be able to recover any data that may have been lost. To do this, all internal fabric connections to the Brick must be disconnected during the recovery; otherwise, the system will detect the recovered drives and proceed with the creation of a new blank LUN.

Even if the data on the Brick is non-essential, you should not accept the alert without first contacting the Pillar World Wide Customer Support Center. Otherwise, the Pillar Axiom system might become inoperable.

9.47 Creating Logical Volumes Immediately After Replacing a Failed Drive

When a failed drive in a SATA Brick storage enclosure is replaced, the system copies the data temporarily stored on the spare drive to the drive replacement. This write operation is called *copyback*. When the copyback operation starts, the unused storage on the affected Brick temporarily becomes unavailable for new allocation. This condition is necessary to guarantee that the unused space is correctly reconditioned.

Approximately once a minute, the Slammer updates the active Pilot control unit with allocation information. If a request for a new allocation happens to arrive during the same minute the copyback operation started, a discrepancy can exist between the allocation information on the Pilot and that on the Slammer. If the discrepancy is large enough to make the difference between success and failure of the allocation, the Pilot can believe the allocation request will succeed, even though the Slammer will fail the request. If this situation occurs, the allocation request appears to fail for no reason because it seems enough free space exists.

If the same request is re-submitted a minute later, after the Pilot has received the next update from the Slammer, the request will correctly fail because the Pilot now has the information that the free space produced by the copyback operation is not available. After the copyback operation has made enough progress in reconditioning enough free space, that same allocation request will succeed, as it would have if there had been no copyback operation.

9.48 System Configuration Vulnerability in Single-SATA Brick Systems

When a factory-fresh, single-SATA Brick system is first powered up, the system configuration database (which is the Pillar Axiom data store containing many system settings) will be configured on that Brick. This database will be doubly-redundant with both instances residing on the same Brick (on two data LUNs). If additional Bricks are added later, the system configuration database will not be migrated to locate the different copies on separate Bricks for higher protection against single-Brick failure.

Important! The system configuration database is still vulnerable to that original, single SATA Brick going offline, as will be user data. If this database goes offline, the Pillar Axiom system attempts immediate emergency shutdown and restart operations to restore access to the system configuration. If this action fails, the system shuts down due to the configuration database being unavailable.

As of release 5.0, through the use of Storage Domains, you can cause the system to migrate the redundant instances of the system configuration database onto separate Bricks by performing the following general actions:

1. Add new Bricks to the system.
2. Place them into a new Storage Domain.
3. Set the new Storage Domain as the primary domain.

The Pillar Axiom system will then migrate the configuration database onto the Bricks in the newly primary domain and stripe the database across multiple Bricks.

Tip: For assistance in determining which Bricks contain the system configuration database, contact the Pillar World Wide Customer Support Center.

9.49 Replacing a Slammer Motherboard Can Cause Several Status Changes

When replacing a Slammer motherboard tray, while the new tray is inserted and powered on, the GUI may initially show the new motherboard status as green (Normal), indicating that the motherboard is functionally OK.

While the Pillar Axiom system attempts to place the new motherboard in service, it will check the Slammer PROM version to see if it matches the installed software version. If necessary, the system updates the PROM to match the current software package version. If the update occurs, the GUI may change the status of the new motherboard to red, because the FRU is offline during the PROM upgrade process, which takes a few minutes. If the upgrade completes successfully, the GUI shows the status of the new motherboard as green and restores it to service if configured to do so.

Important! Do not power cycle a replacement motherboard in the middle of a Guided Maintenance operation; otherwise, the motherboard could be rendered inoperable if it is in the process of updating the Slammer PROM.

As the new motherboard is brought online, the Pillar Axiom system attempts to perform a failback operation on that motherboard. The system checks the PROM version during the failback sequence and, if necessary, updates the PROM. When PROM update completes, the system resets the Slammer CU, which causes the CU to go offline and then go through the failover-failback sequence again. This double failover-failback sequence is normal for Slammer motherboard replacements, if the revision currently installed on the Pillar Axiom system is higher than that on the replacement motherboard.

Important! Do not power cycle a Slammer CU that is in a Failback Pending or Failback in Progress state.

If the preceding operation succeeds, the Slammer CU status becomes Normal and is brought online automatically for SAN Slammers.

Important! Do not attempt the Verify function during Guided Maintenance of a Slammer motherboard. The verification will fail.

Important! Do not run Slammer diagnostics without explicit instructions and assistance from the Pillar World Wide Customer Support Center.

9.50 Testing Failure Recovery from Loss of Slammer Power

Testing failure recovery by removing all power from a Slammer in a dual-Slammer system may result in the remaining Slammer going offline or the system restarting.

The power inputs, power supplies, management paths, and control units in the Slammer are redundant, making this failure injection a multiple failure. Perform this type of multiple fault injection only when it is acceptable to lose the services of the remaining Slammer. Consider contacting Pillar Professional Services who can assist in testing Slammer failover through the use of a support-level CLI command.

9.51 Uploading Empty Files Through the GUI Is Not Allowed

Uploading of zero-byte (empty) configuration or other system files using the Pillar Axiom Storage Services Manager can cause system errors. This restriction applies to all system interfaces where a file upload is allowed.

9.52 DHCP Behavior on the Pilot

In the current release, the DHCP feature has the following behavior characteristics:

- Dynamically assigns only the public IP address of the Pilot.
- Locks the two private IP addresses.
- Retains DHCP settings during a Pilot failover.
- If the IP address is updated through `axiomcli`, the updated address and the status of the DHCP setting are not reflected in the GUI until the Pilot restarts or fails over.
- Updates the values correctly without a need to restart when you change back to static addresses.

Note: You should configure the two private IP addresses to be on the same network as the dynamically assigned public IP address; otherwise, the private interfaces may not work.

If DHCP is enabled on the Pilot and DNS lookup is available on the management console, you can log in to the Pillar Axiom system using the system name rather than its IP address.

9.53 Changing Slammer Port IP Addresses from Static to Dynamic

When changing the IP address from static to DHCP, the change is not instantaneous. The static IP address is retained until a lease is obtained and the system refreshes the status. In other words, the GUI will experience a delay in reporting the newly acquired DHCP address when you change a port from static IP to DHCP.

9.54 VSS Provider Event Numbers Incorrectly Mapped to Descriptions

For VSS Provider events sent to the Windows event log, the VSS Provider plug-in doesn't correctly set the mapping of event number to event description. However, when you click on the event to see its properties, the event text is viewable.

9.55 Disabled/Excluded Slammer States

Repeated failure of a Slammer control unit (CU) can result in that CU becoming non-operational. If the repeated failures occur during normal operation of the Pillar Axiom system, the system marks the CU as Disabled. Power cycling a Slammer CU multiple times can trigger a Disabled state.

If any Slammer CU failure (including a warmstart) occurs during startup, the system marks the CU as Excluded. This behavior can be triggered by power cycling the Slammer CU just once during startup.

If the system completes startup successfully, the system attempts to recover each Excluded Slammer CU:

- After startup completes, the CU should transition to Failed Over.
- If the CU is part of a SAN Slammer, the system should then attempt to transition the CU to Failback as long as the buddy CU on that same Slammer is online. If the Failback succeeds, the system puts the CU Online. If the Failback fails, the system repeats the failover-failback sequence until the Slammer CU failure threshold (*see Section 9.56 below*) is reached. If that threshold is reached, the CU will be Disabled.
- If the CU is not detected, but the other CU is active, a missing CU may show a status of Failed Over when it is really Offline, as in powered down or not connected to the private management network.

Note: Contact the Pillar World Wide Customer Support Center for assistance in recovery of Slammer CUs that have been marked Disabled.

9.56 Slammer Warmstart and Startup Failure Handling

Slammer control units (CUs) have independently maintained fault thresholds. If any of these are exceeded, the system will disable the Slammer CU to allow the rest of the system to continue operation:

- If a Slammer CU warmstarts four times in one hour, it will fail over. If there is a successful failback, the warmstart history count will be cleared.
- If a Slammer CU fails three times in one hour or four times in one week, it will be disabled.
- If a Slammer CU fails during the startup process, it will be Excluded from the startup. If the system startup succeeds, the system will attempt to recover the CU with the failover/failback process. If that fails, the CU will be disabled.

Contact the Pillar World Wide Customer Support Center for recovery assistance for any Slammer CU that is Excluded or Disabled.

9.57 Hardware Lockout Due To Repeated Power Cycling

Important! Do not repeatedly power cycle the control unit (CU) of a Pillar Axiom Slammer. Doing so might automatically trigger the hardware-fault lockout mechanism, which would result in the CU being disabled.

The current thresholds for repeated power cycles are:

- Two power cycles in a 1-hr period
- Three power cycles in a 24-hr period

If either of these thresholds is exceeded, the affected hardware component may be locked out (Disabled).

Note: The above thresholds and Disabled status apply to repeated failover-failback sequences as well.

Contact the Pillar World Wide Customer Support Center for assistance in recovering and restoring the components to service.

9.58 Powering Off a Pillar Axiom System

If you expect to shut down the system for longer than 12 hours, you should remove the batteries from the Slammer after you shut down and power off the system. Reinstall the batteries before restarting the system.

CAUTION! Make sure the system has been placed in Shutdown status before powering it down or removing the batteries; otherwise data loss may result.

Tip: The Read Only system alert should indicate that the Pillar Axiom system is Read Only because the system has been shutdown. If that alert indicates the system is Read Only because shutdown failed, contact the Pillar World Wide Customer Support Center for assistance.

9.59 Battery Removal

When replacing a Slammer battery on a running Pillar Axiom system, be sure to use Guided Maintenance. After you click **Prepare System** in the GUI, Guided Maintenance prepares the system for replacement of the battery:

- Flushes cached data to the Bricks.
- Places all the logical volumes on the target Slammer control unit in Conservative mode.
- Powers down the battery charger.

After the system is prepared, Guided Maintenance displays a completion message and enables the **Next** button. At that point, you can safely remove the battery.

9.60 Battery Insertion

After the insertion of a battery into a Slammer control unit, the battery will show a Warning status in the GUI for a period of time. How long the Warning status remains depends on the charge level of the battery. The time can be up to 18 hrs for a severely discharged battery. If the battery takes longer than 18 hrs to reach a full charge, you should replace the battery. Contact the Pillar World Wide Customer Support Center for assistance in checking the state of the batteries or for a replacement.

Tip: After a battery replacement, wait at least two hours and then collect a set of logs, including the Slammer logs. The Support Center can use those logs to determine whether your battery is charging properly.

9.61 When Pinned Data Is Not Written to Stable Storage

Pinned data is the data stored in Slammer cache in the event of the failure of both control units or one of the arrays; this data cannot be written to stable storage. If the conditions are resolved but the pinned data is not written to stable storage, contact the Pillar World Wide Customer Support Center.

9.62 Hardware Component States

The state of Slammers, Bricks, and the Pillar Axiom system may not update correctly if the Pilot receives hardware events in rapid succession. To view these hardware states, wait 15 seconds. If the Pillar Axiom Storage Services Manager does not show updated states correctly, refresh your browser display.

9.63 Vulnerability Scanners May Report False Positives

The Pilot management controller is protected by means of a firewall to help prevent unauthorized access. Some vulnerability scanners may report a false positive by claiming a large quantity of UDP ports are open when in fact they are not.

9.64 OpenSSH Has Some Vulnerabilities

Pillar addresses security vulnerability in these ways:

- Pillar Axiom systems disable Secure Shell (SSH) access by default.
- Pillar Axiom systems recognize only certain listening addresses for the ports.
- Pillar Axiom systems use shell programs that exist in the cgi-bin directory associated with the web server only to bring up the main GUI login page and to identify whether the login is secure.
- Pillar Axiom systems do not install the source.asp file available for Apache web servers.
- Pillar Axiom systems do not use the Active Server Pages (ASP) feature that runs under mod_perl for Apache servers.
- The Apache modules that Pillar Axiom systems do support are as follows:
    - alias_module

  - auth_basic_module
  - authn_default_module
  - authn_file_module
  - authz_default_module
  - authz_groupfile_module
  - authz_host_module
  - authz_user_module
  - cgid_module
  - core_module
  - dir_module
  - env_module
  - filter_module
  - http_module
  - mime_module
  - mpm_worker_module
  - rewrite_module
  - setenvif_module
  - so_module
  - ssl_module
- Pillar Axiom systems will support the latest version of OpenSSH and SSL in an upcoming release.

9.65 Shutting Down a Pillar Axiom System

In some cases, when you attempt to shut down the system, the system may not shut down but instead return an error message because of a critical task that is running. Before attempting a shutdown or restart, ensure that no background tasks are running. If you are unsure about a task or are unable to cancel a task, contact the Pillar World Wide Customer Support Center for assistance.

9.66 Login to Oracle EM Plug-In Fails

The login to the Oracle EM plug-in may fail even though you enter the correct username and password. This is fixed in Oracle EM release 10.2.0.3. When using earlier releases, place a blank character at the end of the password field, which will enable the login to work correctly.

9.67 SAN Host Associations and LUN Mappings

9.67.1 Definitions

Initiator. A Fibre Channel (FC) port name (WWN) or an iSCSI initiator name (IQN).

Host. An entry in the Pillar Axiom Storage Services Manager GUI or Pillar Axiom CLI that represents a group of one or more initiators. A host is typically used to group all the initiators in a SAN host system. When a LUN is mapped to a host, the Pillar Axiom system makes the LUNs visible to all the initiators owned by that host.

Unassociated host. A host that is created automatically by the Pillar Axiom system when an initiator is discovered on the SAN. It contains a single initiator and has the same name as the initiator.

Associated host. A host that is created by the administrator using the GUI or CLI commands.

APM host. A host that is created when Pillar Axiom Path manager (APM), which is running on a SAN host system, starts communicating with the Pillar Axiom Pilot. An APM host includes all the initiators that APM discovered on the SAN host system. The administrator cannot modify the list of initiators owned by an APM host.

The Pillar Axiom system always creates an unassociated host for each discovered initiator. Initiators may subsequently be moved into associated hosts by the administrator or into APM hosts by APM. An administrator can map LUNs to all three host types.

9.67.2 Managing Unassociated and Associated Hosts

The administrator can use the GUI or the CLI to move initiators from unassociated hosts into associated hosts, and between associated hosts.

9.67.2.1 Moving an Initiator From an Unassociated Host

When the administrator moves an initiator from an unassociated host to an associated host, the LUN mappings of the unassociated host are added to the associated host and the unassociated host is deleted. The original mappings belonging to the associated host are unaffected. If a mapping from the unassociated host uses a different LUN number for a LUN that is already mapped to the new host, or if a mapping from the unassociated host uses a LUN number that is already in use for a different LUN in the new host, the mapping from the unassociated host is deleted and a LUNMappingDeleted event is generated.

9.67.2.2 Moving an Initiator From an Associated Host

When the administrator moves an initiator from one associated host to another, the mappings of the previous associated host are removed from the initiator, and the initiator is mapped with the mappings of the new associated host. The previous associated host (even if it has no initiators left) is left with its LUN mappings in place. No mappings associated with the previous associated host are moved to the new associated host. The previous host can be manually deleted, if it's no longer needed.

9.67.2.3 Deleting a Host

When deleting an unassociated or associated host, a new unassociated host is created for each initiator that is connected. When deleting a host, you can remove all LUN mappings from the initiators or keep all of the LUN mappings belonging to the deleted host with each new unassociated host. No unassociated host is created for a disconnected initiator.

9.67.2.4 Renaming a Host

Unassociated hosts and associated hosts cannot be renamed. You can achieve a similar effect by creating a new associated host with the desired name and transferring the initiators from the existing host to the new associated host.

To effectively rename an associated host while preserving the mappings, ensure that the old associated host's initiators are shown as connected, and then delete the old associated host using the option to preserve mappings. The Pillar Axiom system creates a new unassociated host for each of the initiators, with the same mappings as the original associated host. You can then transfer the initiators from each of these unassociated hosts into the new associated host.

9.67.3 Managing APM Hosts

The name of an APM host, and the list of initiators that make up the host, are managed by APM running on the SAN host. For APM to be able to manage host entries in the Pillar Axiom system, APM must be able to log in to the Pillar Axiom Pilot. For information on how to ensure that APM can log in to the Pilot, see the Installation Guides for your versions of APM.

When APM on a SAN host logs in to the Pillar Axiom Pilot, APM sends a message that gives the name of the host and the list of initiators that it found on the host. If an APM host with this name does not currently exist in the Pillar Axiom system, the host is created and the identified initiators are transferred into it. If an APM host with this name already exists, but its list of initiators is not the same as that now reported by APM, initiators are transferred in or out of the existing definition to match those in the new message.

After APM has logged in and the APM host entry has been created, APM is reported as "Communicating" in the Pillar Axiom Storage Services Manager for as long as APM stays logged in (see the **Configure** > **Storage** > **Hosts** display screen). APM sends a similar message from time to time (for example, if the configuration changes at the host) for as long as APM is in a "Communicating" state. If necessary, the host definition is updated in the Pillar Axiom system each time a message is received.

9.67.3.1 Moving an Initiator From an Unassociated or Associated Host

When an initiator belonging to an unassociated host or an associated host is moved into an APM host, the LUN mappings of the old host are added to the APM host. If the old host is an unassociated host, it is deleted. If the old host is an associated host, its LUN mappings stay in place even if no initiators are now associated with it. The old associated host may be manually deleted if it's no longer needed.

The original mappings belonging to the APM host are unaffected. If a mapping from the old host uses a different LUN number for a LUN that is already mapped to the APM host, or if a mapping from the old host uses a LUN number that is already in use for a different LUN in the APM host, the mapping from the old host is not added and a LUNMappingDeleted event is generated.

9.67.3.2 Moving an Initiator From Another APM Host

When an initiator belonging to another APM host is moved into an APM host, the LUN mappings of the old host are not added to the new host. The old mappings to the initiator are removed and the initiator is mapped to the configuration of the new owning APM host.

9.67.3.3 Combinations

If a message from an APM host requires initiators to be moved from more than one old host, initiators and mappings are added from unassociated hosts to the APM host first, followed by adding initiators and mappings from the associated hosts, followed by adding initiators from APM hosts. The order in which hosts of the same type are dealt with is not specified.

9.67.3.4 Removing an Initiator From an APM Host

If a message from a SAN host running APM does not include an initiator that is currently included in the APM host, the initiator is removed from APM host and all mappings are removed from the initiator. If the initiator is connected to the Slammer, an unassociated host is created for that initiator.

9.67.3.5 Deleting an APM Host

When deleting an APM host, a new unassociated host is created for each initiator that is connected. When deleting a host, you can remove all LUN mappings from the initiators or keep all of the LUN mappings belonging to the deleted host with each new unassociated host. No unassociated host is created for a disconnected initiator.

Note: If APM is running on the SAN host and is able to discover and log in to the Pillar Axiom Pilot, it might recreate the APM host entry at any time.

9.67.3.6 Renaming an APM Host

APM hosts always have the hostname reported in the message from APM that is running on the SAN host. If you want to rename a SAN host that is running APM, use the following sequence of steps to create an APM host with the new name and which has the same mappings as the old APM host.

1. Stop the APM service or daemon on the SAN host. See the Installation Guide for the version of APM being used for information on managing the service or daemon.
2. Ensure that at least one of the host's initiators is connected.
3. Delete the old APM host.
   Be sure to select the option to preserve the mappings.
4. Use the appropriate operating system procedures to rename the SAN host.
5. Start the APM service or daemon on the SAN host.

APM then logs in to the Pillar Axiom Pilot and sends a message reporting the new SAN host name. The Pillar Axiom system then creates an APM host for the new name and transfers in all the initiators and their LUN mappings.

9.68 Host Mappings Can Disappear Under Certain Circumstances

If a SAN host with associated initiators was defined in the Pillar Axiom system before release 2.5 and that host subsequently has initiators added and removed in a single operation (either using the GUI or the PerformAssociateInitiatorsToHost CLI request), the defined SAN host and the associated initiators may disappear.

To recover the lost SAN host and its initiators, re-associate the initiators to the SAN Host again.

Tip: When modifying an associated host created in releases prior to release 2.5, follow these steps to remove and add initiators:

1. Remove the unwanted initiators.
2. Click **OK**.
3. Add the desired initiators.
4. Click **OK**.

9.69 Overloaded Systems Can Result in Multiple System Warmstarts

The Pillar Axiom Pilot software is multi-threaded, and the number of threads available for handling system events and user-requested operations is limited. When a large amount of system activity requires Pilot intervention, this maximum number of threads can become fully utilized, leaving no threads available for user-requested operations.

A large number of requests for action can result in slow responses and the active Pilot control unit (CU) could fail over to the passive CU. Should this situation arise, wait until the number of system events requiring Pilot action is reduced, freeing up the Pilot for additional user-requested operations.

9.70 IPStor Fails To Recognize LUNs After a Full System Upgrade

IPStor systems do not recognize Pillar Axiom LUNs after a system upgrade from a Pillar Axiom 500 to a Pillar Axiom 600. IPStor systems recognize only the Pillar Axiom 500 as a valid target. To resolve this situation, request from your IPStor vendor the code update necessary for compatibility with the full Pillar Axiom product line.

Pillar Axiom 500 systems can be upgraded to the Pillar Axiom 600 platform. However, as you do so, Bytes 16-31 of the Inquiry Data response (the "ASCII Product Identification" field) for each LUN will change from "Axiom 500" to "Axiom 600". This change may cause some SAN appliances, such as IPStor, and possibly some hosts to not recognize the LUNs unless the appliance or host is reconfigured.

9.71 Getting Accurate SNMP Filesystem Performance Metrics

If an SNMP request is made within 60 seconds of the last request of the same type, the system returns the same information as it did in the prior request. To ensure you get the latest information, wait more than 60 seconds between identical SNMP requests.

10 Technical Documentation Errata

The following sections describe topics in the technical documentation that were not able to be corrected in time for the current release of the Pillar Axiom 500 and 600 system software.

10.1 *APM 3.1 Installation Guide and Release Notes for Oracle VM Server 2.2*

On page 31, Step 5 should read as follows:

> *For each package (APM or Multipath Tools), click the name of the package to download.*

10.2 *APM 3.0 Installation Guide and Release Notes for RHEL4*

- At the top of page 13, the following sentence does *not* apply to APM 3.0 for RHEL4:

  > *You will need to rename "blacklist" to "devnode_blacklist" in the multipath.conf file if you want to blacklist internal SCSI devices.*

  The above sentence applies only to the APM 3.0 for CentOS4 and OEL4 platforms. APM 3.0 for RHEL4 users should ignore this instruction.

- On page 17, the Operating Limits tables is missing the following entry:

Connect to LUNs	Maximum = 256 visible from each Pillar Axiom system

- On page 24, the QLogic HBA driver version number `8.01.04` that is listed in the path is incorrect. The correct driver version number is `8.02.14.01`. The correct path should be `cd qla2xxx-8.02.14.01`.

10.3 *APM 3.0 Documentation for RHEL5.2, OEL5.2, and CentOS 5.2*

The iSCSI Initiator configuration instructions that appear in the following Pillar Axiom Path Manager books are incorrect:

- *AxiomONE Path Manager 3.0 Installation Guide and Release Notes for RHEL 5.2*
- *AxiomONE Path Manager 3.0 Installation Guide and Release Notes for OEL 5.2*
- *AxiomONE Path Manager 3.0 Installation Guide and Release Notes for CentOS 5.2*

In particular, the incorrect instructions appear in the following two sections of those books:

- "Configure the iSCSI Software Initiator"
- "Start the iSCSI Software Initiator Service"

If you are configuring the iSCSI Initiator for the RHEL5.2, OEL5.2, or CentOS5.2 version of APM 3.0, follow the instructions in the README file for the iSCSI Initiator that comes with your Linux distribution and ignore the iSCSI Initiator configuration instructions in the APM documentation.

If you have questions or require detailed iSCSI Initiator configuration instructions for your installation, please contact the Pillar World Wide Customer Support Center.