Pillar Axiom® 600

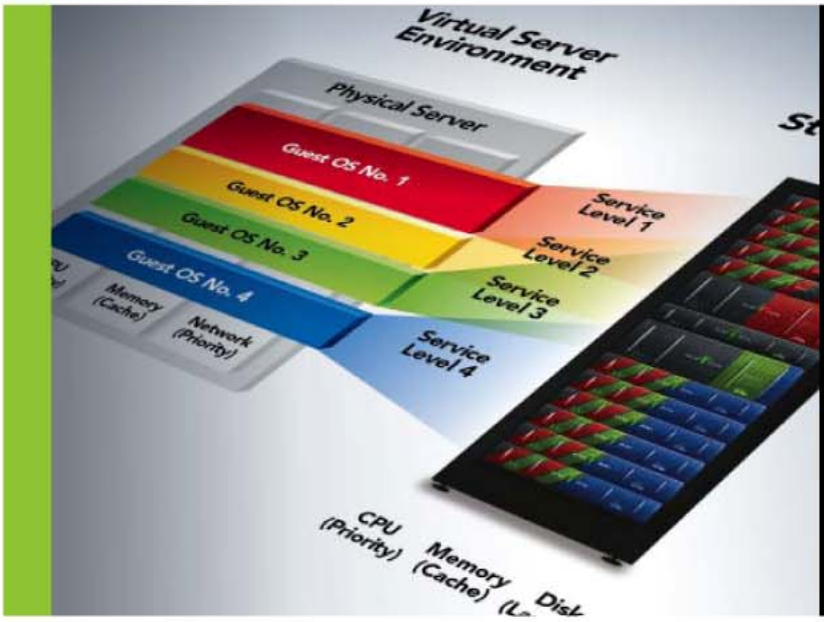

Customer Release Notes

For Release 3.4

pillar DATA SYSTEMS

Pillar Axiom® 600	Document Number: 4420-00026-2800
	Document Title: Customer Release Notes

Revision History

Rev Description	Rev Date	Effective Date
Release 3.4 GA	2009-07-22	2009-08-28
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
Release 02.08.00	2008-01-10	2008-01-14
Release 02.07.00	2007-10-29	2007-11-05 [expand list of fixed issues]
Release 02.07.00	2007-10-01	2007-10-22
Release 02.06.00 GA	2007-06-29	2007-06-29
Release 02.06.00 Beta	2007-06-08	2007-06-08
Release 02.05.00	2007-04-13	2007-04-13
Release 02.04.00	2007-02-11	2007-02-12

1 Purpose

This document describes new features, capacities, configuration requirements, operating constraints, known issues and their workarounds, and other items for release 3.4 of the Pillar Axiom 600 storage system. 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 Data Systems authorized representative.

2 Product Release Information

Release 3.4 is a maintenance and feature enhancement release of the unified NAS/SAN Pillar Axiom 600 storage system.

2.1 System Enhancements

This software update provides several feature and performance enhancements for NAS and SAN Pillar Axiom storage systems.

With release 3.4, solid state disk (SSD) storage arrays are now supported on Pillar Axiom 600 systems.[†] An SSD Brick storage enclosure provides the following features:

- Provides 650 GB of raw storage dedicated to the Premium QoS bands only. [‡]
- Provides faster performance compared to Serial ATA (SATA) and Fibre Channel (FC) Bricks.
- Consumes approximately 125 W of power, compared to around 300 W for FC Bricks.
- Up to four SSD Bricks for each Slammer storage controller can be added to existing SATA-only configurations.[§]

Note: For Brick configuration details, see Section 5.2.

Tip: For cabling information, contact the Pillar World Wide Customer Support Center.

Important! Adding SSD Bricks to an existing system is disruptive to the data path.

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 3.4.

Note: Beginning with release 3.0, the Pillar Axiom system dropped the Triple Redundancy Quality of Service (QoS) option for *new* volumes (see also Section 2.4.3). However, using Double Redundancy in a Pooled RAID 10 array effectively provides quadruple redundancy, if that is desired.

For the list of defects that this release resolves, see Table 8 beginning on page 29.

2.2 Changes to the Back Up to Disk Feature

Beginning with release 3.0, full-block backups (Volume Backup) of logical volumes (filesystems and LUNs) are no longer supported in the AxiomONE Storage Services Manager (GUI). Because of that, administrators need to be aware of the following when updating a pre-3.0 system to release 3.4:

- Existing volume backups will automatically activate and become full-fledged LUNs and filesystems. If a volume already exists with the same name, the name of the activated backup will contain a numeric suffix to uniquely identify it. LUN backups that are activated will be mapped but not to any specific hosts.
- To create volume backups, use the new backup-to-disk feature, Clone FS and Clone LUN, both of which allows for the creation of inactive clones.

[†] For this release, SSD Bricks are not supported on storage configurations containing Fibre Channel (FC) Bricks.

[‡] SSD performance bands do not support thinly provisioned LUNs and filesystems.

[§] SSD Bricks must be added at the head of a Brick string and connect directly to a Slammer containing version 2 private interconnect modules (PIMs).

In earlier releases, administrators could choose "Backup to Disk" to create a full-block copy (having a QoS setting of Archive) for an existing logical volume. Because the 3.1 and higher releases support partial-block snapshots of logical volumes, "Backup to Disk" has been removed from the GUI. Full-block snapshots with the ability to set QoS levels and data path accessibility, however, are still available through the Copy action in the GUI. Furthermore, the command line interface (CLI) continues to provide the volume backup facility through the pdscli utility.

Beginning with the 3.1 release, all backups created by means of the CLI (`pdscli`) are partial snapshots, not full-block copies set to Archive QoS. The CLI continues to support `GetAllVolumeBackups` and associated requests. The GUI displays these as part of the Clone FS and Clone LUN lists.

Table 1 Data replica equivalencies between the CLI and the GUI

CLI		GUI data replica
Data replica	**Command**	
Volume backups	CreateVolumeBackup DeleteVolumeBackup GetAllVolumeBackups GetVolumeBackupDetails RestoreFromVolumeBackup	Inactive clones
Clones or snapshots	GetAllSnapLUNHierarchies GetAllSnapLUNs GetSnapLUNDetails CreateSnapLUN ModifySnapLUN DeleteSnapLUN PerformRebuildSnapLUNRedundancy PrepareSnapLUN SyncSnapLUN CreateCloneFS DeleteCloneFS GetAllCloneFSs GetCloneFSDetails ModifyCloneFS GetAllCloneFSHierarchies PerformRebuildCloneFSRedundancy RestoreFileSystemFromCloneFS	Active clones

Note: The CLI commands `CreateLUNCopy` and `CreateFileSystemCopy` have been retained for backward compatibility. With this release, they also create inactive clones.

2.3 Available Licenses

Feature	Comments	Pillar Axiom model		
		300	500	600
CIFS and NFS Protocols	NAS Slammer: One protocol included at no cost.	•	NA	NA
	NAS Slammer: Optional.	NA	•	•
Fibre Channel and iSCSI Protocols	SAN Slammer: One protocol included at no cost.	•	•	•
SNMP	Automatically included.	•	•	•
NDMP	Automatically included.	•	•	•
Pooled RAID 10	Automatically included.	•	•	•
SecureWORMfs	Optional.	•	•	•
Snap FS	Automatically included for NAS Slammers.	•	•	•
Copy Services bundle	Optional: – NAS: Clone FS + Volume Copy / Backup – SAN: Clone LUN + Volume Copy / Backup	•	NA	•
Clone FS	Optional. For Pillar Axiom 300, see above.	NA	•	•
Clone LUN	Optional. For Pillar Axiom 300, see above.	NA	•	•
Volume Copy / Backup	Optional.	NA	•	•
Thin Provisioning: FS	Optional.	•	•	•
Thin Provisioning: LUN	Optional.	•	•	•
Volume Copy	Optional.	NA	•	•
Volume Backup / Restore	Optional.	•	•	•
Support Tool	Optional.	•	•	•
VTL Appliance	Optional.	•	•	•

The features that you have licensed each have a unique feature key. These feature keys are listed in a separate document** and are associated with your Pillar Axiom storage system. These keys allow access to the features that you have licensed. Please keep the license key document in a secure place for future reference. For additional information, please contact Pillar Data Systems at 1-877-4PILLAR or go to www.pillardata.com.

** Feature keys are listed on the packing slip associated with your product shipment. You can also obtain the keys on the Support portal.

2.4 Pillar Axiom Software Update

The 3.4 release may be installed through the GUI Software Update process. This update will be disruptive to the data path.

If you prefer, you can have Pillar Professional Services update your system software. For more information, refer to Table 4 Contact information.

2.4.1 Prerequisites to Updating Pillar Axiom Software

Before a Pillar Axiom system can be updated to release 3.4, the system must be running release 3.3 (or higher).

Important! For *single*-Slammer Pillar Axiom 600 systems, Pillar highly recommends that you provide additional intra-control unit (CU) cabling that takes full advantage of the Opteron CPU in both Slammer controllers. Contact Pillar Worldwide Customer Support Center for assistance in cross cabling the two CUs. If you don't provide this extra cabling, after upgrading to release 3.4, you will receive an Administrator Action reminding you to provide that cabling.

Important! Contact the Pillar World Wide Customer Support Center for assistance in updating any system below release 3.3. Do not attempt to perform this software update without this assistance.

2.4.2 The Update Process

Pillar recommends that you perform all updates during a maintenance window.

Note: Because this update is disruptive, all I/O and workload on the Pillar Axiom system must be quiesced.

2.4.2.1 Before Starting the Update

Before you start the software update process to release 3.4:

- Ensure no background processes are running.
  You can check for running background processes from the system GUI by pressing **Display All Background Processes Running** in the lower right corner of each GUI screen.
- Before starting the update, stop all user applications.

  Important! Updating a system to release 3.4 is disruptive to the data path.

It is not necessary to disconnect channel attachments from the Pillar Axiom system.

2.4.2.2 Starting the Update

When you start the update process, even though it is possible to select individual software modules, *you should select all modules* that have selection boxes available; otherwise, the update may not proceed due to a potential incompatible combination of software modules.

Important! Be sure to scroll through the entire list so you can select all available modules.

CAUTION: In some cases, the system will allow the updating of individual software modules, even though doing so will create a fatal system error. In particular, attempting to update Pilot Software to release 3.4 from release 3.3 without updating the Slammer Software will create a fatal system error. If this situation should occur, contact the Pillar World Wide Customer Support Center for assistance.

Important! Do not initiate any system or storage actions of any kind until the update process completes.

2.4.2.3 During the Update

If the Pillar Axiom system contains one or more NAS Slammers:

- CIFS connections are dropped (however, most applications will attempt to reconnect automatically). Dropping CIFS connections is not a timing issue and is not related to how the update process is performed. Connections are dropped because the system does not maintain connection states across failover or warmstart events.
- NFS access will not survive the disruptive update process.
- NDMP protocol does not tolerate server warmstarts. Therefore, when this disruptive update happens in the middle of a backup session, all backup applications will fail.

Important! After updating the system software to release 3.4, you cannot reverse the update and downgrade to an earlier release level without explicit action by the Pillar World Wide Customer Support Center.

After a successful update, all system components will report Normal in the GUI and user data will be available.

2.4.3 System Free and Total Capacity May Be Reported Differently in the CLI

At the *system* level, the GUI and the CLI will continue to report the same amount of total and free capacity after upgrading from earlier 3.x releases to 3.4. However, at the *priority band* level, the CLI (by means of the `GetStorageConfigDetails` request) may report different values for the total and free capacities after the upgrade. This difference at the band level occurs because the system determines the total capacity in a band by subtracting 50 GB from the total system capacity before computing the per-band values. If the free capacity in a band changes, the total capacity will change as well.

2.4.4 Triple Redundancy Will Cause an Alert

After upgrading to release 3.4, if any logical volumes have a triple redundancy QoS setting, the system will generate a Call-Home event.

Note: Pillar World Wide Customer Support will notify you if such an alert is received by Pillar.

This alert occurs to help customers prepare for future releases in which Triple redundancy will no longer be supported.

2.4.5 Software Versions for This Release

Software provided in this release includes the versions listed in the following table.

Table 2 Pillar Axiom Software versions

Software module	Software version
Pilot OS	03.04.00
Pilot Software	03.04.00
Slammer PROM	03.01.00
Slammer Software	03.04.00
Brick Firmware	07.06.06
Brick SATA2 Firmware[††]	00.06.06
Fibre Channel Brick Firmware	00.66.06
Brick Disk Drive Firmware	*Version depends on the disk drive capacity and whether it is the spare or array disk drive.*

After updating your system, check your software versions in the GUI by going to Support > Software Modules.

Note: If your Pillar Axiom system does not have FC Bricks installed, the GUI will not display the Fibre Channel Brick Firmware entry.

Important! Starting with the 2.x releases, the WWN was changed to use a common base World Wide Node Name. When updating a Pillar Axiom system from 1.7.x to 3.4, be prepared for a change in how WWNs are managed. For more information, see Section 9.18 on page 45.

2.5 AxiomONE Path Manager Software

The AxiomONE Path Manager (APM) software provides for the following:

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

The current release of APM for SAN hosts varies by platform, as shown in Table 3.

Table 3 APM support

Operating system	OS release	APM release	Platforms
AIX	5.2	2.1	All platforms
	5.3	2.2	All platforms
Community Enterprise Operating System (CentOS)	4.5 and 4.6	3.0	32-bit x86, 64-bit x86, Itanium
	4.7	3.1	32-bit x86, 64-bit x86
	5.2	3.0	32-bit x86, 64-bit x86

[††] Applies to both SSD and version 2 SATA Bricks.

Operating system	OS release	APM release	Platforms
HP-UX	11i v1 and v2	2.0	All 64-bit platforms
	11i v3	2.0	All 64-bit platforms
Novell SUSE Linux Enterprise Server (SLES)	9 SP2 and SP3	2.2	32-bit x86, 64-bit x86, Itanium
	10 SP2	3.0	32-bit x86, 64-bit x86
Oracle Enterprise Linux (OEL)	4 u4	2.2	32-bit x86, 64-bit x86
	4 u5 and u6	3.0	32-bit x86, 64-bit x86
	4.7	3.1	32-bit x86, 64-bit x86
	5 u2	3.0	32-bit x86, 64-bit x86
Red Hat Enterprise Linux (RHEL)	3 u7 and u8	2.2	32-bit x86, 64-bit x86, Itanium
	4 u3 and u4	2.2	32-bit x86, 64-bit x86, Itanium
	4 u5 and u6	2.3	POWER
	4 u5 and u6	3.0	32-bit x86, 64-bit x86, Itanium
	4.7	3.1	32-bit x86, 64-bit x86, POWER
	5.0	2.0	32-bit x86, 64-bit x86, Itanium
	5.2	3.0	32-bit x86, 64-bit x86
Solaris	8	01.04.05	SPARC with non-Sun HBAs
	9	01.04.05	SPARC with non-Sun HBAs
	9	2.0	SPARC with Sun HBAs
	10	2.0	SPARC, 64-bit x86
Windows	2000 Server	2.3	32-bit x86, 64-bit x86, Itanium
	Server 2003 Server 2008	3.1	32-bit x86, 64-bit x86, Itanium

For complete details concerning APM requirements, refer to the *AxiomONE 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.

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

3 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.

4 Support

Pillar Data Systems provides various levels of customer service on a contract basis. If you have purchased a service contract from Pillar Data Systems, authorized Pillar Data Systems personnel will perform support and repair according to the terms and conditions of that agreement.

Table 4 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: 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 storage systems. Hardware that does not conform to Pillar specifications or is not a Pillar-supplied part voids the warranty and may compromise data integrity.

4.2 Access to Pillar Axiom Systems

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

- The AxiomONE Storage Services Manager (GUI)
- The AxiomONE 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 support.pillardata.com.
2. Click the Log In link at the top right of the navigation bar.
3. Enter your username and password.
4. Click the Login button.
5. On the main view, go to the Downloads area.
6. Click the appropriate software release. Download details for that software appears on the right side of the display.
7. Click the green download button.

Note: To obtain a license to enable a feature that is in addition to those that you initially purchased, contact a Pillar account representative. (See Table 4 Contact information.)

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 a CIFS environment, see the *Pillar Axiom Windows Integration Guide for NAS Systems*.

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 storage system and how to configure them:

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

The above documents are available on the **Support > Documents** page in the GUI and on the Pillar Data Systems Web portal at www.pillardata.com/techdocs/.

5 Pillar Axiom System Limits

This version of the Pillar Axiom storage 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—File Servers, filesystems, LUNs, shares, exports, snapshots, 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 *Administrator's Guide* PDF file (search for *Ranges for Field Definitions*).

Table 5 Operating limits of a Pillar Axiom NAS systems

Item	Description and Range
File Servers	Maximum = 8 for each Slammer ‡‡ §§
VIFs for each File Server	Minimum = 1 Maximum = 32
VIFs for each Slammer port	Maximum = 16*** (any of which may belong to any File Server)
VLANs for each File Server	Minimum = 0 Maximum = 32
Static and default network routes for each File Server	Minimum = 0 Maximum = 8 default, 16 static
NIS alternative file size	Up to 50 MB
Volume groups	From 1 to 5000, out to four levels
Filesystems	Minimum = 1 Maximum (including Clone FSs) = 1024†††
Filesystem size	Minimum = Greater of: • 1% of maximum size • 1 to 2 GB (depending on the QoS) Maximum = system capacity
Directory quotas	Unlimited

‡‡ In multi-Slammer systems, a File Server's virtual interfaces (VIFs) can be configured on multiple Slammers. The presence of VIFs is what counts against the above limit. Such a File Server is considered to be present on each Slammer on which it has VIFs.

§§ VLAN tagging does not need to be enabled for more than one File Server. Also, as of release 3.0, File Servers no longer require a unique VLAN tag.

*** If VLAN tagging is enabled, the maximum is 5 VIFs for each filesystem.

††† Filesystem maximum limit applies to both the system and to a NAS Slammer.

Snap FSs	Maximum = • 250 for each filesystem • 16,000 for each Pillar Axiom 600 system
Volume Backups	Maximum = • 1024 for each Pillar Axiom 600 system • 1024 for each filesystem Maximum concurrent = 5 Maximum outstanding I/O requests = 128 across all backups
Concurrent NDMP backup and restore sessions	Maximum = 5 for a given system
NFS exports	Maximum = 1000 for each File Server
NFS host entries	Maximum = 4000 for each File Server
CIFS shares	Maximum = 128 for each File Server
User security groups	Maximum = 1024 for each CIFS user
CIFS connections	Maximum for each Slammer (combined memory for both control units): • 1200 for 12 GB combined memory • 3000 for 24 GB combined memory

Table 6 Operating limits of a Pillar Axiom SAN systems

Item	Description and Range
SAN LUNs	Maximum = • 2048 visible for each system • 256 visible for each host • 1024 visible for each SAN Slammer
SAN LUN size	Minimum = Greater of: • 1% of maximum size • 1 to 2 GB (depending on the QoS) Maximum = system capacity

Volume Copies (full block snapshots)	Maximum = • 128 for each Pillar Axiom 600 system • 12 active for each LUN
Clone LUNs (partial block snapshots)	Maximum = number of unallocated SAN LUNs, up to 1024
iSCSI	Maximum = • 256 TCP connections for each iSCSI port • 256 iSCSI Initiators for each iSCSI port • 32 persistent reservation registration keys for each LUN • 512 simultaneous commands for each iSCSI port

5.2 Slammer and Brick Configuration Limits

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

Number of Slammers	Minimum number of Bricks		Maximum number of Bricks
	Supported	Recommended	
1[‡‡‡]	2	3	32
2	4	All NAS or all SAN Slammers = 5 NAS / SAN combo Slammers = 6	64
3	6	8	64
4	8	8	64

The first two Brick storage enclosures in all configurations must be the same type and capacity.

Note: Bricks containing solid state disks (SSDs) are not supported in storage configurations containing Fibre Channel (FC) Bricks.

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

[‡‡‡] For single-Slammer Pillar Axiom 600 configurations, the minimum number of Bricks is two of the same type. However, for mixed configurations, the minimum number of Bricks is three: 2 SATA + 1 SSD or 2 SSD + 1 SATA.

6 System Requirements

6.1 Using Browsers on Windows XP Operating Systems

When using Microsoft Windows XP, set the Windows Desktop appearance to something other than the default XP theme to ensure that lines and boxes are displayed correctly.

Important! If you use the default theme, some controls (such as radio buttons) will appear as though they were not there.

Configure the browser:

- Set security to medium-low (or lower) to enable the security certificate.
- Enable image support (if not enabled).
- Enable JavaScript.
- For Microsoft Internet Explorer, disable the Script Debugger.
- Set the displayed text size to the smallest comfortable viewing size.

When logging into the AxiomONE Storage Services Manager using Secure HTTP, you may 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.

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.
- The external switch ports must be set to auto-negotiation for the Pilot interfaces.
- 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 AxiomONE Path Manager communicates with the Pilot over secure, encrypted XML. If the Path Manager is installed on a SAN host, that host will require an Ethernet interface for communication with the AxiomONE 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

NAS data paths require gigabit Ethernet connections. Both fiber and copper are supported.

SAN data paths require 1 Gbps, 2 Gbps, or 4 Gbps 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 NDMP Requirements

For a list of data management applications (DMAs) that Pillar Axiom systems support, see the latest *Pillar Axiom Support and Interoperability Guide*.

6.3.1 File Server

The NDMP subsystem uses the networking configuration from a single File Server, which can be selected by means of the AxiomONE Storage Services Manager (GUI). Currently, a File Server must be set in the NDMP configuration portion of the GUI.

6.3.2 NDMP Command Interface

The NDMP command and response interface on a Pillar Axiom system is the Pilot management interface. Data movement is performed over the data path interfaces on the Slammer storage controllers. Be sure that any external NDMP backup servers are able to reach the Pilot management IP addresses.

6.3.3 Virtual Interface (VIF)

Only one VIF is required. For local backups, the networking configuration must be on the Slammer control unit (CU) to which the tapes are attached. The tape menu in the GUI lists the CU as a control unit number.

There are two ways to ensure there is networking on the CU with tapes:

- The first method is to create the File Server on the CU with tapes (or alternatively move it once it has been created).
- The second method is to create a second VIF on the CU to which the tapes are attached.

  Note: The File Server used must be the File Server listed in the NDMP configuration.

6.3.4 Fibre Channel (FC) Tape Library LUNs

Even though the AxiomONE Storage Services Manager (GUI) allows you to configure tape library LUNs from 0 to 255, the FC tape driver only supports eight (0-7). To avoid difficulties, don’t define library LUNs above number 7.

6.4 Power-Off Requirements

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

Note: If you will be powering down the system for more than a day, remove the Slammer batteries so they do not discharge.

6.5 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 which case, 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.

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 4 Contact information).

Table 7 Known Pillar Axiom server issues

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
ALL	In extremely rare situations when restarting a Pillar Axiom system, errant configuration information may prevent the system from restarting.	Contact Pillar World Wide Customer Support This issue will be fixed in a future release.
	Following an upgrade from release 3.1 to 3.3, a filesystem that was offline (due to a geomap startup bug) failed to come online following the upgrade that was intended to correct the original offline failure.	A warm start of the Slammer control unit or a restart of the Pillar Axiom system will clear the condition. This issue will be fixed in a future release.
	An internal data structure to map IDs appears to have been corrupted for unknown reasons and returned invalid data. That caused an abort when the data was de-referenced, but since the structure was in volatile memory, it was rebuilt on the restart, and the system returned to normal function.	This issue will be fixed in a future release.
	Currently the system does not explicitly alert an administrator to the possible presence of loopback connectors installed on the system Fibre Channel (FC) ports.	Check the system for the installation of loopback plugs and remove any of these found before operating the system. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	The system may mis-detect tape devices after a warmstart or restart. If the tape device returns an error, the Pillar Axiom tape driver may not retry the error correctly.	Rescan for the tape devices using the CLI or the GUI. This issue will be fixed in a future release.
	Because releases earlier than 3.2 do not understand the software package for version 2 Bricks, when you upgrade from release 3.1 to 3.3 (and higher), an Administrator Action is generated with the error: “Inconsistency of SW dependencies.”	Ignore the message. Continue with the update in the normal way. This issue will be fixed in a future release.
	The Pilot uses a version of OpenSSH (OpenBSD Secure Shell) that has several known vulnerabilities.	Although there is no workaround, Pillar Axiom systems have numerous security features that make it unlikely anyone could break in. *(See Section 9.85 for more information.)* An upgrade to the latest release of OpenSSH will be available in a future release.
	Additional backend FC loop recovery code meant to detect and resolve loop issues was introduced in the upgraded code. Unfortunately, manufacturing test loopback plugs were accidentally left in the ports, and these have the same loop issue signature the new recovery code looks for, and will trigger the recovery action. The problem is that the recovery code does not give up as long as the "issue" remains.	Check and make sure there are no loopback plugs before you upgrade. This issue will be fixed in a future release.
	If a system software update is initiated while Pilot tasks are in progress (such as SecureWORMfs scans and log collection), the update will not begin and no information is provided that the update is waiting for the tasks to complete.	• 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 (or, if desired, cancel them). This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	While a logical volume (LUN or filesystem) is being copied, if an attempt is made to modify the source volume and the modification causes the volume to be migrated, the attempt will fail. Such modifications include changing the relative priority or RAID level of the source volume.	Do not modify a logical volume that is the target of a copy operation. This issue will be fixed in a future release.
	A faulty Enclosure Services module in a Fibre Channel Expansion Brick can prevent the Expansion Brick and its RAID partner from coming on line.	Replace the faulty ES module. 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 disk 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 disk drive issue separately. This issue will be fixed in a future release.
	Sometimes, when a Fibre Channel cross connection between two Slammer control units is missing, the Administrator Action misreports the missing connection.	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 proxy server 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.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	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.
	Sometimes during Pilot failover, if the startup time for the newly active Pilot control unit (CU) is lengthy, the system generates a hardware compatibility Administrator Action, which is not cleared when the failed Pilot CU restarts and begins communicating with the active Pilot CU.	Remove and reattach a Brick bezel, which should cause the system to refresh the hardware compatibility Administrator Actions. This issue will be fixed in a future release.
	If a logical volume is being restored from a backup, attempts to modify the attributes of the volume (such as the Slammer control unit assigned to the volume) will fail.	Wait for the background restore to complete before modifying volume attributes, or modify the attributes before running the restore operation. This issue will be fixed in a future release.
	The Verify Storage Redundancy operation does not report any progress.	Leave the system undisturbed until the operation completes. This operation could take hours or much longer. This issue will be fixed in a future release.
	In rare cases, the system status (exhibited for example on the Health page or in Administrator Actions) may not be entirely accurate.	Depends on the component. Contact the Pillar World Wide Customer Support Center. This issue will be fixed in a future release.
	It is possible to assign volumes (such as filesystems and LUNs) to volume groups such that the total allocation size of the volumes exceeds the maximum defined capacity of the volume group.	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
	Very large system information or Call-Home files cannot be downloaded to a client system using the web download function. You get an error popup: A system error has occurred: Cannot download file. Try again or contact the Pillar World Wide Customer Support Center.	Try using the transfer function. Alternatively, reduce the amount of information in the bundle by selecting "most recent" or by breaking the information into two bundles: one with all of the Bricks and another with all of the Slammers. This issue will be fixed in a future release.
	Trying to modify any characteristic of a scheduled software update can lead to the system error: Cannot add fully qualified name (FQN) to table because the name already exists Enter a different FQN.	Delete the scheduled update you wish to change and create a new one with the desired characteristics. This issue will be fixed in a future release.
	If a disk 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 disk drive is used automatically.	This issue will be fixed in a future release.
	The PDSCLI command "GetStorageConfigDetails" returns raw capacities based on a RAID 5 geometry. This does not account for capacity requirements when using Pooled RAID 10. The RAID 10 raw capacity is 1/2 of the listed capacities returned by this command.	When using the PDSCLI command "GetStorageConfigDetails" on a Pooled RAID 10 system, divide the listed capacities by 2 to get the actual capacity. This issue will be fixed in a future release.
	When changing the Pilot management IP from static to DHCP in the GUI, the change does not get committed.	Run the CLI command ModifyManagementConfig with DHCPEnabled set to true and the pilot1 and pilot2 IPAddress fields set with proper static values. This issue will be fixed in a future release.
NAS	After a share is disabled using the GUI, a new CIFS connection can be established to this share after the Pillar Axiom system is restarted, even though the share had been disabled earlier (A “disable share” checkbox was ignored by the restart configuration routine).	To permanently delete this configuration, delete the share. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	When an ACL is changed, the "change" time is not updated to the current time. One consequence of not updating the change time is that an NFS client can still believe that it can safely use its cache. For example, run "ls" on a directory that is accessible to everyone. Now change the ACL of the directory to have more restrictive permissions. When the "ls" command is repeated, the NFS client gets the attributes of the directory and checks to see if the change time has been updated. If it hasn't, then the NFS client concludes that the directory hasn't changed and it can rely on its cached data. So, the repeated "ls" also works when it should be failing.	Invoke a write operation, such as creating a dummy subdirectory, which would modify the “change” time and cause the NFS client to flush its cache and honor the new ACL. This issue will be fixed in a future release.
	When a file or directory is created by a CIFS client and then accessed by a NFS client, subsequent NFS operations do not update the access control list (ACL) or the mode correctly: NFS permissions appear to be more restricted than intended.	Ignore the NFS mode bit displays. Legal access should still be permitted. This issue will be fixed in a future release.
	A mixture of CIFS and NFS operations can result in the creation of an incorrect ACL. For example, sometimes a new directory can result in an incorrect mode of 0000.	After chown or chmod, use CIFS to reapply the desirable ACLs. This issue will be fixed in a future release.
	If a filesystem is offline, filesystem snapshot schedules cannot be deleted. A “Delete filesystem snapshot task failed” Administrator Action will be generated.	Do not put the filesystem offline until after you have deleted the snapshot schedule. If the filesystem is already offline, contact the Pillar World Wide Support Center. This issue will be fixed in a future release.
	When a Clone FS is created from a non-WORM filesystem and the clone is then converted to a SecureWORMfs filesystem, a later attempt to delete the parent filesystem (and all its children) through the GUI can cause problems for other objects.	Don’t attempt to delete the non-WORM parent filesystem from the GUI. Note: Pillar recommends that you not create SecureWORMfs clones of non-WORM filesystems. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	For a SecureWORMfs filesystem, the Protected File Integrity Scan start time may be off by an hour from the time listed in the GUI if there has been an intervening change to or from Daylight Savings Time.	This issue will be fixed in a future release.
	While a SecureWORMfs scan is in progress, operations on that WORM filesystem (such as cloning) that require the filesystem to be quiesced may cancel the scan.	Wait for the scan to complete before performing those operations, or restart the scan. This issue will be fixed in a future release.
	Events are logged when a *user-initiated* daily protection scan of a SecureWORMfs filesystem completes but not when a *scheduled* daily protection scan completes.	This issue will be fixed in a future release
	A Pillar Axiom system may have trouble joining a CIFS domain when one or more of the CIFS domain controllers are offline.	Join the domain when all of the domain controllers are available. This issue will be fixed in a future release.
	When using Computer Management or Microsoft Management Console to change the owner of a share, if the user selects the option to "Replace owner on subcontainers and objects", the request fails.	Change the ownership in this way: i. Change the owner of the share, without the "Replace owner on subcontainers and objects" option. ii. Repeat the step to change the owner of the share (to the same owner) but this time, do select the "Replace owner on subcontainers and objects" option. This issue will be fixed in a future release.
	When copying a filesystem on a single-Slammer system, system performance may decrease.	A decrease in performance is expected on single-Slammer systems. As such, you should schedule filesystem copy operations during off-peak hours, particularly when a Clone FS is the source of the copy request. This issue will be fixed in a future release.
	The TCP/IP statistics GUI page under Health > Performance > NAS Protocols always shows link aggregation (LA) status as disabled, even when LA is properly configured.	Use the System > Networking page to verify LA status. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	For a Slammer that has 12 GB memory for each control unit, when the number of CIFS connections reaches approximately 1300, new connections are rejected until some existing connections are closed.	Try either or both of the following: • Close unused applications that have open connections to the CIFS server. • Unmap shares that are no longer required. This issue will be fixed in a future release.
	If both control units (CUs) in a NAS Slammer are failed over and the system is already running with another NAS Slammer, you cannot fail back the CUs in the failed over NAS Slammer.	Place the system in shutdown mode and restart the system. This issue will be fixed in a future release.
	A filesystem may switch from battery-backed journaling mode to conservative mode if it runs out of battery-backed memory (BBM) journals. This may happen when existing in-use journals cannot be flushed out in time, which can be indicative of a heavy load, a sluggish backend, or a code defect that causes a journal leak. Once the filesystem has switched to conservative mode, it stays in that mode until the Slammer control unit warmstarts.	This issue will be fixed in a future release.
	If memory consumption approaches the limit (around 400 KB) on a particular CIFS connection, individual CIFS requests on that connection will fail. In rare conditions, the Slammer control unit may drop a CIFS connection and possibly warmstart.	This issue will be fixed in a future release.
	Joining a domain may not succeed if TCP traffic to port 464 is blocked. Port 464 is used by the KPASSWD protocol to set the domain account password for a domain client. The KPASSWD protocol supports both UDP and TCP. However, a Pillar Axiom system fails the operation if the TCP port is blocked.	Unblock port 464 for TCP network traffic. Ensure that the domain controller and all routers in the path have port 464 unblocked for TCP traffic. This issue will be fixed in a future release.
	If a Slammer warmstarts after a request to create a Snap FS has been issued, in very rare circumstances the request may fail.	Reissue the Snap FS request. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Under some circumstances, an opportunistic lock (OpLock) is not released on a file. This causes a software fault when another lock request is made.	This issue will be fixed in a future release.
	Under certain conditions, the Pillar Axiom system will fail to recognize user credentials, such as when a client requests to map a drive. Such requests will be rejected.	Contact the Pillar World Wide Customer Support Center for assistance. This issue will be fixed in a future release.
	Under certain conditions, when a CIFS client writes data to the Pillar Axiom system, the write might fail, the CIFS connection is terminated, and the log contains this record: Got EBADF error, terminating CIFS connection	Contact the Pillar World Wide Customer Support Center for assistance. This issue will be fixed in a future release.
	If the TCP ports for LDAP services on a Domain Controller are blocked, a request by a File Server to join the domain will fail.	Port blocking and failover from TCP to UDP when TCP is blocked applies to Kerberos and kpasswd but not to LDAP. The domain controller should not block any port for LDAP services. This issue will be fixed in a future release.
	When a CIFS server has joined a domain in ADS mode and a user successfully authenticates using ADS and connects to that server, if the administrator then assigns CIFS local group and the user subsequently tries to make use of the privileges as part of CIFS local group assignment, such as trying to access a folder for which he or she did not previously have permissions, the attempt will fail.	Purge the Kerberos ticket on the client (or have the user log off on the client) or retry the operation. This issue will be fixed in a future release.
	Mounting a filesystem on a secondary virtual interface (VIF) will fail if the connection uses UDP through a firewall.	Use TCP. If UDP must be used, limit the VIFs for the File Server to one. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Un-checking the “Require Authentication” option on the File Server CIFS tab in the GUI does not allow a client with an anonymous connection to the File Server to obtain a list of shares with either "net view" or "net use".	Leave the “Require Authentication” option checked. This issue will be fixed in a future release.
	Because of a network issue, the request from the Pillar Axiom system generated no response, eventually causing the system to fail a health check after the health check timed out.	Check the network and Domain Controller if this happens repeatedly. This issue will be fixed in a future release.
	When using a device with less than 1 Gbps link speed to connect to a NAS Slammer, the I/O Port Details screen will indicate 0 Gbps for the Negotiated Link Speed.	Use the Filesystems or TCP/IP statistics pages to view the actual speed in Mbps. This issue will be fixed in a future release.
	In rare circumstances, a request to modify a filesystem can fail if a system warmstart occurs while the request is executing.	Retry the failed request. This issue will be fixed in a future release.
	When there is high volume CIFS traffic and multiple accesses to the same file with opportunistic lock (OpLock) enabled, the OpLock will not be released in a timely manner leading to a health check timeout.	This issue will be fixed in a future release.
	When both a directory and one of its subdirectories are being shared, restrictive permissions on the top directory may prevent some users from mapping the subdirectory. In particular, if the top directory has permissions that prevent some users from traversing this directory, these users will be unable to use either share.	Allow everyone the right to traverse all directories leading up to any directory that is being shared. Being able to traverse a directory does not imply that these users will be able read or modify the contents of that directory. This issue will be fixed in a future release.
SAN	Due to an initial failure, where the maximum number of LUNs were consumed, a secondary warmstart occurred.	This issue will be fixed in a future release.
	Due to a very small timing window in buffer handling when a Slammer warmstart occurs, additional warmstarts may occur.	This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	When you specify automatic assignment of LUNs to Slammer control units and then delete some LUNs in a particular priority band, the system may not properly balance the remaining LUNs across the owning control units for that band.	When you create a LUN, manually assign it to a specific Slammer control unit. This issue will be fixed in a future release.
	Certain fatal and some ECC errors in the SAN network interface module (NIM) can disrupt heartbeats between the two Slammer control units (CUs), causing CU failover.	Replace the problem 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.
	When using the Capacity Planner to create a very large quantity of LUNs, it fails and the GUI shows an error dialog with no text, just an OK button.	Try creating 20 LUNs at a time. If you want to create more than 20 LUNs, try using a CLI script. This issue will be resolved 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.
	Sometimes when mapping a LUN, the mapping fails with error 11202 but, internally, the system makes it appear to the host that the mapping succeeded. The host will see a LUN as that LUN number mapped to it, even though the mapping is not really mapped to the host. Also, the invalid LUN mapping shows on the GUI pages that display mappings.	Restart the system to clear the out-of-sync mapping in the SAN and to re-synchronize the Pillar Axiom system and the SAN. This issue will be resolved in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
iSCSI	If an APM host uses iSCSI to connect to a Pillar Axiom system, and it uses an iSCSI Initiator Name which is the same as its host name, then the entry for that host in the AxiomONE Storage Services Manager will be continually deleted and recreated. The host entry will disappear and reappear intermittently in the GUI.	Ensure that the iSCSI Initiator Names configured on hosts that use iSCSI to connect to a Pillar Axiom system are different from all hostnames used by APM hosts connected to that Pillar Axiom system. The iSCSI Standards require that iSCSI names follow particular formats, as specified in RFC 3720 (http://tools.ietf.org/html/rfc3720#section-3.2.6). If hosts are configured to use iSCSI Initiator Names which conform to these requirements, it is extremely unlikely that they will be the same as any host name. This issue will be fixed 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 start-up, 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.

Table 8 Resolved Pillar Axiom server issues

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
3.4	ALL	Due to a change in the way that mutli-control unit processing of logical volumes was implemented in 3.3, the possibility exists that on a non-disruptive upgrade from 3.2 to 3.3.4 or earlier may result in logical volumes coming up in the offline state.
		When one control unit (CU) fails over and is followed by a CU warmstart, the warmstart may timeout and fail due to an internal processing deadlock.
		If an administrator enters 0 for both soft and hard quota limits in the GUI, they will get an error message indicating that both fields cannot be zero.
		A race condition exists in the 64-bit timer handling routine of the Pillar Axiom software that under very rare circumstances can cause a control unit to warmstart. The clock_handler() ISR is being preempted by other interrupts in the system.
		In 3.x software prior to release 3.3.5, if a Slammer control unit (CU) fails over and the surviving CU warmstarts immediately (before the failover completes), write operations may occur in write-through mode only, resulting in poor write performance.
		If RAID controller 1 in an FC Brick has previously had a warmstart due to an error, it is possible that a rebuild that is initiated by the Pilot will get stuck.
		A bug exists in the Axiom software related to handling a temperature related automatic shutdown. In this case, a bad fan FRU caused a Slammer control unit (CU) to detect the over-heating and to fail over to it buddy CU. The surviving CU flush out orphaned I/O destined for the failed CU. This bug causes the clearing of the orphaned I/O to be slow. The system interpreted the slowness as a hang and warm started the surviving CU to recover the perceived hang.
		The Lost Data flag could not be cleared for a non-standard logical volume, such as a clone of a filesystem or LUN.
		In software versions prior to 3.2, if a large number of write I/O transactions were in process when a control unit (CU) failover occurs, a warmstart may result due to an excessive delay in timing out mirrored writes to the failed CU.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		After a disk drive failure, it is possible for the firmware to encounter an exception as the rebuild to the spare is started. If this should occur, the rebuild to the spare will not complete and the system will be unable to perform a copyback operation.
		When a Brick becomes inaccessible, warmstarts may result due to the inability of the system software to access critical operating metadata.
		It is possible for a failing Fibre Channel disk drive to get into a state where it drops off and comes back on the backend loop repeatedly. This behavior causes a disturbance in the backend loop, which can greatly hinder the performance of the Brick.
		This problem was due to a defect in the code for managing data migrations. When adjusting internal data structures to reflect the migration, the code made a reference to memory that had already been reclaimed. Often this happens soon enough after de-allocation that there is no problem. In this case, the memory reference caused an infinite recursion stack overflow, causing a Slammer control unit to warmstart. That warm start actually fixed the problem, because the memory contents were then regenerated correctly.
		A slow memory leak was detected that could cause a warmstart of a Slammer control unit after one year of uptime.
		Following the update of the system software to R3.3, a pointer was not properly initialized resulting in an improper memory access.
		A component responsible for re-configuration during Slammer control unit (CU) failback after a failure sets a "failback in progress" state that it expects to clear when failback configuration is done, and it can shift back to normal operation. That state was not saved over a warmstart. So, in this case, when another CU happened to warm start during the failback, the state was lost, and the system behaved as if it were restored to normal state prematurely.
		A component responsible for re-configuration during Slammer control unit failback after a failure sets a "failback in progress" state that it expects to clear when failback configuration is done, and it can shift back to normal operation. In addition to the state, it also sets a watchdog timer with a limit (30 minutes) for the time in that state. It assumed that if it exceeded 30 minutes, it was probably a bug in clearing the state. In this defect, however, re-configuration had to do large amounts of disk I/O, and it legitimately took longer than the watchdog allowed, causing the component to abort configuration.
		A defect was detected in the low level kernel routines of the Pillar Axiom system software that can lead to a failover of a control unit. This is a very rare occurrence that has to do with a segmentation fault while the kernel is in the middle of sending a pulse to a kernel process.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		Under very rare circumstances, an error may be encountered during system shutdown. The GUI reports the status of the system shutdown operation as "Shutdown Failed".
		If you change the name of a LUN or filesystem to a name that already exists, the system may erroneously allow the change, which may prevent the administrator from making subsequent modifications to that LUN or filesystem.
		A system warmstart may result when fewer than half of the configured Bricks are online, which can cause a Slammer control unit to become disabled.
		When a number of Bricks have gone offline, the system may warm start.
		When viewing the simulation results in the Capacity Planner in the GUI, the summary values for Used system capacity, Total system capacity, and Remaining system capacity are inaccurate and should be ignored.
	NAS	In very rare conditions, the pathnames or file name may contain illegal characters. In those cases, the Pillar Axiom system returns an illegal sequence error causing the CIFS server to warm start.
		When a Snap FS is accessed during failback, a small window can occur when the snapshot is mounted for access but the active filesystem ownership has not been reestablished. This situation causes the NAS control unit (CU) to warmstart repeatedly leading to the CU being disabled.
		Under a race condition, a CIFS process can use a TCP socket that has been freed causing a NAS Slammer control unit to warmstart.
		When a file is being deleted and another file with CIFS reservation is closed at the same time, the machine may sometimes warm start.
		In very rare conditions, the path names or file name may contain illegal characters. In those cases, the Pillar Axiom system should have returned an error code indicating this is an invalid pathname. Instead, it returned illegal sequence error causing the CIFS server to warm start.
		The filesystem was not handling a case where the linked list was empty. As a result, while de-referencing an empty linked list, a segmentation fault was encountered.
		In a rare condition, the filesystem database tracking user opportunistic locks (OpLocks) can get out of sync. In those cases, a stale lock record could cause the NAS control unit to time out on a health check due to the wait time of breaking OpLock that would exceed the desirable time.
		In rare cases, the OpLock (opportunistic lock) records in the database are not cleaned up properly causing OpLock related issues.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		A NULL quota tree was getting associated because of a variable not being set, this NULL quota tree was causing the segmentation fault. The code has been rectified to handle the NULL quota tree condition.
		Validation scans on SecureWORMfs filesystems can slow system performance significantly.
		A SecureWORMfs audit directory (`/.audit`) and its logs are not readable and cannot be made readable in Windows.
		If more than one host entry exists for a given NFS export, the order of processing the entries could be different from the one shown in the GUI display (usually after inserting host entries in the middle of the list).
		When a CIFS client has been given an opportunistic lock (OpLock) that allows writes to be delayed and the connection between the client and the CIFS server is experiencing network difficulties, the client can experience delayed-write errors.
		For an extremely large configuration having more than 1000 filesystems, a Slammer warmstart sometimes causes control unit failover.
		Both file deletion and truncation to a smaller size are asynchronous operations. Multiple file deletions and truncations can be batched and processed in the background. When a batch includes large files, the delete or truncate will take too many CPU cycles and lead to system warm start. This is especially true if all files to be deleted or truncated are in the same directory.
		While using the supporttool wbinfo to enumerate local groups, the Pillar Axiom system might warmstart under certain circumstances.
	SAN	When host systems that are attached to the Pillar Axiom server drive I/O at a higher rate than the server can support, the server returns BUSY. The fix now returns TASK_SET_FULL instead of BUSY.
		On systems running VMware ESX the host system may not properly select preferred IO paths resulting in non-optimized accesses. This issue is resolved in Pillar Axiom release 3.3.8 and above.
		When an administrator changes the selection of a host in the mapping tab during a LUN create or modify request in the GUI, the administrator can create a map based on the available numbers before the list of available numbers is updated, resulting in a number that is not available.
		Some operating systems may respond better to a TASK_SET_FULL I/O return as opposed to a BUSY return in situations where the Pillar Axiom system is being overdriven by I/O.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		In rare situations the Slammer may warmstart in situations where I/O is overdriven to the Slammer.
		Following the update of the system software to release 3.3, a pointer was not properly initialized resulting in an improper memory access.
		In situations where a mis-configured system results in a large number of non-optimized access I/Os, the potential exists for a system warmstart.
		If, in the process of configuring SAN Host mappings for LUNs during failback of a Slammer control unit, MCC encounters a missing SAN Host software record, no mappings for any LUNs beyond that point will be configured.
		If a warmstart occurs following deletion of Clone LUNs, the data structures for these Clone LUNs are not properly reclaimed, which may result in running out of available Clone LUNs and then another system warmstart.
		Due to a problem handling quorum disk handling clustering failover may not successfully occur.
		If two or more APM hosts connect to a Pillar Axiom management interface using the same source IP address (as seen by the Pilot), then the Pilot will repeatedly run out of resources. This causes the Pilot to dump core and restart, briefly interrupting any management connections. The APM hosts will appear to continuously connect and disconnect from the Pilot. Configurations of this sort can be created by using Network Address Translators (NATs) between the hosts and the Pilot; although the hosts themselves may be configured with unique IP addresses, the connections at the Pilot may all appear to come from the IP address used by the NAT.
		In the unlikely situation where clustered hosts are being used without benefit of APM and non-optimized paths result, if a RESERVE command is sent from the host to the Pillar Axiom system and the non-optimized accesses result in a delayed execution whereby the host aborts the RESERVE command at the same time the Pillar Axiom system has processed the RESERVE command, a warmstart may occur.
	iSCSI	In iSCSI environments with noisy power, high temperatures, or very heavy I/O loads, a system warmstart may result.
		Downgrade of iSCSI Initiator from `Initiator-2.08-build3825-x64fre.exe` to `Initiator-2.06-build3497-x64chk.exe` may cause a problem with iSCSI session login handling.
		In rare situations the iSCSI chip used in the Axiom may hang in certain network error conditions resulting in an Axiom warmstart and an access outage of up to 5 minutes.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		It is not possible to create multiple iSCSI sessions to Pillar Axiom systems through a QLogic 4052C iSCSI HBA when using the latest QLogic 4052C HBA software and the latest Microsoft iSCSI Initiator software.
3.3	ALL	After upgrading a multi-Slammer system to release 3.2, when a write operation to a double redundant logical volume fails or a similar issue occurs that causes the mirrors to lose synchronization, the system may panic and warm start repeatedly.
		In release 3.x, live Brick logs cannot be collected as part of doing a manual system log collection.
		On rare occasions, after a Slammer control unit (CU) fails over, the fail back may complete incorrectly, leaving the CU in a Normal state but with the IP address of one of the virtual interfaces being inactive.
		Under extremely rare circumstances, a Slammer control unit (CU) will experience multiple, consecutive software failures due to incorrect internal (non-data path) memory handling during normal operation. This will usually lead to the CU being disabled for excessive errors.
		When upgrading a Pillar Axiom system from 2.x to a 3.x release earlier than 3.2.7, the internal growth increment that is selected for a large LUN or filesystem may be too great, which results in the volume using system capacity inefficiently (by wasting space) or not being able to grow at all (by exceeding available system space).
		An NDMP backup operation sometimes can hang because of an internal snapshot issue. When this occurs, NDMP will cause the Slammer to warm start but the backup operation will continue to hang.
		After upgrading from release 2.x to release 3.2, attempts to modify the system Global Settings might fail.
		Under some conditions, if a Slammer Fibre Channel cross connection fails, the system fails to generate an Administrator Action that identifies the missing connection.
		If a system is in a Startup Failed state, attempting to shut down or restart the system from the GUI or CLI may result in filesystem corruption.
		The Doho (Qatar) and Riyadh (Saudi Arabia) time zone setting sets the time to GMT-3 instead of GMT+3.
		During Guided Maintenance, if a Slammer battery replacement fails diagnostics, the battery will not be charged, eventually causing the Slammer control unit to go into Conservative mode.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		If the passive Pilot control unit (CU) is turned off or fails during a software update process, that Pilot CU may not have the target versions of the Pilot OS and Pilot Software installed in the update. Also, the system may not indicate that a failure has occurred.
	NAS	Creating and deleting a large number of Snap FSs over a lengthy period of time without an intervening restart or warmstart can cause an out-of-memory event and a subsequent warmstart.
		When a filesystem is homed on (assigned to) one Slammer control unit (CU) and the user accesses the filesystem through the buddy CU, sometimes that access can cause the Slammer to warm start.
		Reclaiming the space formerly allocated to large files and reporting the reclaimed space as free can take a long time.
		When an administrator attempts to have a Pillar Axiom system join a Window 2008 domain, the request will fail because Windows 2008 DC is not supported.
		When several Snap FSs are created on the same filesystem and the mounting of the first snapshot fails, the mounting of the second snapshot may also fail.
		When the user accesses a Snap FS while the source filesystem is in journal recovery mode, the Slammer can repeatedly warm start, even after the filesystem has recovered and is in a Normal state.
		Reassigning a filesystem to a different owning Slammer control unit, particularly while I/O is in progress to that filesystem, can cause system crashes and possible data loss or corruption.
		NLM locking operations can sometimes hang, which after a few minutes causes the system to warm start. Typically, the NFS client automatically retries the lock operation.
		In rare circumstances, disk usage reports might show less free space than anticipated after big files or a large number of files are deleted or truncated.
		In release 3.2, deleting large files or a large number of files can take longer to reclaim data blocks than in earlier releases. When queried (for example, with `df -k`), the amount of free space does not immediately appear.
		If Clone FSs are actively used for I/O during a failover-failback event, the clone may go offline for a few seconds before coming back online.
		If a user enters a default route for a File Server that already exists in the system, the system responds with an error code that is not handled by the GUI properly, which results in an Error tab with no dialog showing what caused the error.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		The CIFS server scans trusted domains every five minutes. If one of the servers within the trusted domain list fails while the scan process is running or when a Slammer control unit (CU) is attempting to connect to that server, the CU could warmstart.
		If a snapshot is restored after the filesystem FSCK is performed, there is a chance the snapshot can still have the same corruption.
	SAN	When the VTL license is not installed, creating a single SAN LUN with the Capacity Planner will fail. When "Finish" is selected, the Capacity Planner returns a blank Information pop-up with just an OK in it.
		When using an SNMP client to get specific LUN details (such as statistics), the Pilot management software increases the Pilot workload significantly, which can cause a Pilot failover if the system contains 50 LUNs or more.
		Unable to activate Clone LUN without a Volume Copy license.
		Performance degradation of non-optimized path accesses causes extent lock thrashing, resulting in a PANIC and a system warmstart.
		When a Clone LUN other than the oldest for a snapshot tree is deleted, the snapshot repository space that was allocated to the clone is reassigned to the next older clone. While this reassignment occurs, all older clones and the next newer clone (or the source if the newest clone is being deleted) are not accessible. This can take a long time when much data exists that has to be reassigned and may cause problems for applications and command timeouts within the system.
	iSCSI	The VMware iSCSI Hardware Initiator fails during virtual machine (VM) creation if a 2-MB, 4-MB, or 8-MB block size is used to create the VM filesystem on a Pillar Axiom LUN.
		On occasion, an iSCSI client using a QLogic HBA may see the Pillar Axiom system warmstart, which causes a short interruption in I/O.
3.2	ALL	When updating the system software from 2.x[§§§] to 3.x, if the maximum size and current size of a volume repository differ by a factor greater than 100:1, the update fails with excluded Slammers.
		If a hardware compatibility check is performed while a RAID controller is offline, the system reports hardware compatibility Administrator Actions for the offline controller.
		If an FC RAID controller is either physically removed or warmstarts while heavy I/O is in progress, the surviving RAID controller might warmstart as well.

[§§§] 2.x refers to Pillar Axiom 500 systems. The equivalent release for Pillar Axiom 300 systems is 1.x.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		In rare cases, the generation of a system log event can result in a RAID controller warmstart.
		In R3.x builds prior 030200-004800, when a RAID controller fails, the controller number reported is that of the *partner* controller (for example, if RAID controller 0 fails, controller number 1 is reported). This occurs because the problem is reported by the controller that did not fail. In R3.2 builds and later, the controller number of the failed controller is reported.
		In rare circumstances, after a warmstart, a logical volume (filesystem or LUN) may go offline because of lost data.
		Modifying a volume's QoS while the volume is being restored from a clone can cause both tasks to fail, and possible data corruption.
		On occasion, a Pillar Axiom system may warmstart if the I/O load is too heavy given the number of Bricks in the system.
		During a system upgrade, the Pilot currently decides the system type based on the type of the first Slammer control unit (CU) discovered. If no Slammer CUs are discovered, the Pilot defaults to using the Pillar Axiom 500 upgrade images. For Pillar Axiom 600 systems, the upgrade will fail.
		Sometimes the following event might be reported: “FC path to control unit failed” The likelihood of this storage path failure in the backend increases with increases in traffic load.
		In the unlikely event that a failed drive is replaced with another failed drive, more that one drive will be in warning state and the Brick LUN could be taken off line.
		If AxiomONE CLI is installed on Windows in a directory whose full path contains a directory name that contains a space character, the axiom_login command will fail. The error message indicates that the ID or password is incorrect but the source of the error is the space character in the full path.
		The following event can occur during normal operation of a Pillar Axiom system: “Fibre Channel Path to Control Unit Failed”
		On the AxiomONE Storage Service Manager Storage Usage Summary page, Clone LUNs and Clone FSs may be included in the count, after they have been deleted from the system. This situation is most likely to occur when a large number of clones are deleted in a short period of time.
		If a non-disruptive software update is performed, it is possible for the Bricks to end up in an unusable state with RAID firmware at two different levels.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
	NAS	Thinly provisioned filesystems that have a maximum size greater than 2 TB may go offline when the underlying allocation crosses the 2 TB threshold.
		When a user is a member of too many security groups and attempts to map a drive to a Pillar Axiom share, if the CIFS server is configured to use Kerberos authentication, the server will warmstart, causing a brief pause for other users while their CIFS connections are reestablished. The fix rejects the single user's authentication request.
		When the CIFS server attempts to communicate with the Kerberos Key Distribution Center over TCP, sometimes the CIFS server will encounter an error before it can try UDP and will warmstart, causing a pause while the CIFS connections are re-established.
		Sometimes, when a CIFS client operation results in the Pillar Axiom server attempting to retrieve a share level Access Control List (ACL), the system generates a PANIC event and warmstarts, causing a brief pause while the CIFS connections are reestablished.
		The creation of 100,000 directories for which each has a capacity quota generates the following Administrator Action: Create Quota Task Failed.
		Restoring from a Clone FS to a filesystem that has been taken offline causes the restore operation to fail.
		Sometimes, when the CIFS server receives an obsolete CIFS request (such as SMBSearch) from a client, the server will warmstart, causing a pause while the CIFS connections are re-established. (The fix prevents the warmstart but still rejects the request.)
		When a filesystem is corrupt and one or more snapshots are taken (which now contain the corruption) and the filesystem is subsequently repaired by FSCK, the filesystem (which is now consistent) may go Offline without an Administrator Action to FSCK the filesystem.
		When a client uses Kerberos authentication, if the user is a member of more than 254 groups and attempts to map a drive, the Pillar Axiom system panics and warmstarts. (The fix rejects the user authentication request.)
		Sometimes when a Slammer control unit (CU) fails multiple times, an inconsistency can develop in memory and when an attempt is made later to create a new filesystem, that CU can fail again, clearing the inconsistency.
		A join domain operation may fail when domain controllers are not in synchronization.
		When there are more than 1200 open files from a single connection of CIFS client, a memory allocation failure can occur, resulting in a warmstart rather than an error returned to the CIFS client.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		If the user configures NTP by specifying only the second alternate server (not the primary or first alternate servers), the time will not be synchronized between machines, and no errors will be reported.
		In rare circumstances, when a File Server is being modified and the DNS server becomes unresponsive, the Slammer doesn't respond to a command from the Pilot soon enough before the command times out. In such a case, the Slammer control unit (CU) is failed over.
		When deleting a large file, the filesystem may encounter long delays due to resource contention. On rare occasions, the delay is severe, causing the system to warmstart.
		The AxiomONE Storage Services Manager may not display the correct number of Snap FSs in the Storage > Usage Summary page.
		On rare occasions, there may be conditions in which a filesystem is not able to remove a directory and it returns "directory not empty" to the users. This is because there are hidden files under those directories that can not be displayed by the filesystem.
		When CIFS is configured with Account Mapping enabled and with an NIS server, CIFS server attempts to use NIS to map CIFS user names to NFS user names. This mapping occurs during authentication processing while mapping a drive. If the NIS server does not respond promptly, the lower layers of code will set a timer and retry. In certain cases, the retry time will exceed the client's timeout setting, and the client will drop the connection, failing the authentication request.
		When a large number of files are deleted while the system is busy, the filesystem can be tied up by the delete request, which may lead to the healthcheck detector to falsely assume the delete thread is hung and to warmstart the control unit (CU).
		Because file deletion in the system is an asychronous operation, multiple file deletions can be batched and processed in the background. When a batch includes large files, the deletion may require too many CPU cycles and lead to a system warmstart. This is especially true if all files to be deleted are in the same directory.
	SAN	During network interrupts (port logouts), some jobs that are in the process of being aborted (because they were timed out) can cause the Slammer to warmstart.
		An occasional internal configuration problem with thinly provisioned LUNs could result in a warmstart.
		When using Microsoft iSCSI initiator, logging out with a non-empty command queue on the host causes the Pillar Axiom system to warmstart.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		On rare occasions, a Pillar Axiom system might warmstart during a delete LUN operation.
		When the Thin Provisioning license is not installed, attempts to use the Simulator to create a LUN generates an HTTP error and the LUN creation fails.
		When a Clone LUN is expanded in size and the clone's source LUN has also been expanded in size, reads of the Clone LUN may be redirected to the source LUN even if the requested region of the Clone LUN was previously written.
		Device Mapper may have difficulty in booting up if it encounters a LUN with corrupt vtable entries.

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 and/or configuring a Pillar Axiom storage 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 storage 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 Adding iSCSI HBAs While Heavy I/O is Occurring

During the addition of iSCSI HBAs to the Pillar Axiom system, 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 Administrator Action 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

Both Pillar Axiom user interfaces (the GUI and the CLI) 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.

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, Pillar recommends that you assign the LUN to the Slammer control unit (CU) on which you have the mapping set.

9.10 LUNs Created Through Capacity Planner Accessible by All Hosts

When you create a LUN using the Capacity Planning Manager, the LUN is created without port mapping or masking information. To configure port mapping or masking for a LUN that was created through the Capacity Planner:

1. In the Storage>LUN section of the GUI, click the link for the LUN you want to configure.
2. In the LUN Access section, select the “Only selected hosts” option.
3. Click the Mapping tab.
4. Configure the LUN for mapping and port masking as needed.

9.11 System May Rebalance LUNs Without Prompting the Administrator

If you auto-assign LUNs to a Slammer control unit (CU), the system may move those LUNs to another CU when the system starts up (or restarts). The system may take this action to rebalance the system load by QoS bands. This strategy works well if AxiomONE Path Manager (APM) is installed on all client hosts that use the auto-assigned LUNs.

If, however, the client host does not have APM installed, this strategy may cause Non-Optimized Access events. In this case, Pillar recommends that you explicitly assign all LUNs for non-APM clients to a specific Slammer CU.

9.12 Blacklisting Local Disk Drives on RHEL4 Platforms

Device Mapper on RHEL4 U4 platforms may display local SCSI SAS or SATA disk drives along with the FC disk drives as multipathed. Including local disk 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
```

9.13 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 AxiomONE 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.

9.14 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.

Pillar strongly recommends 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.

9.15 Uploading Software Update Packages over Slow Connections

Pillar does not recommend 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 having a speed of 10 Mbit/sec or greater.

9.16 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.

9.17 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 NAS Slammers, the brief protocol disruption is such that most client applications either time out and recover or see a fast reboot of the Pillar Axiom server.

Important! For Pillar Axiom systems that have one or more NAS Slammers, Pillar recommends that you quiesce much of the I/O to the system before performing a non-disruptive update.

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. For example, the Microsoft Cluster Service will initiate retries in 3 sec and, at 7 sec, will begin reconfiguration and recovery that will probably disrupt any applications using the Cluster Service. Microsoft, however, is promising to fix this issue in the upcoming release of the Windows “Longhorn” server (successor to Windows Server 2003).

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 control unit (CU), Slammer Port Number, and Port Type.

9.19 Possible Errors When Running Unsupported Linux Variants

When attempting to run the Command Line Interface (CLI) application on unsupported Linux variants (such as Fedora Core 3 and Core 4 versions of Linux), you may see the following message:

```
pdscli-Linux: error while loading shared libraries:
libstdc++.so.5: cannot open shared object file: No such file or
directory.
```

These and certain other variants of Linux do not include the necessary libraries in their standard installation. Although these are unsupported Linux variants, you may be able to use these versions of Linux by installing the appropriate version of the compat-libstdc++ rpm package.

Note: If you need support for another operating system, contact your Pillar account representative.

9.20 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 AxiomONE 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.

9.21 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 storage system.

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

9.22 Status of Filesystems and LUNs

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 LUNs or filesystems.

A hardware problem will typically cause filesystems and LUNs to go offline or to a degraded state. Because this is not always the case, you should check the state of the filesystems and LUNs or any associated Administrator Actions that may be listed.

9.23 Change in Default Failback Configuration of NAS Slammers

Automatic failback of NAS Slammers is the default configuration beginning with release 2.0 of Pillar Axiom 500 systems (release 1.0 of Pillar Axiom 300 systems). If this is not desired, automatic failback of NAS Slammer control units (CUs) can be disabled in the GUI Global Network settings menu (System>Global Settings>Network).

Automatic failback of SAN Slammer CUs is always enabled.

9.24 Changing the Time on a Pillar Axiom System

A Network Time Protocol (NTP) server is required for some environments, such as those using CIFS, SecureWORMfs, and some NFS environments. 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 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.

9.25 Automatic Snapshots

When you create a filesystem, the default action is to also create a snapshot schedule that will create filesystem snapshots every four hours. If this schedule is not appropriate for your data, perform one of these actions:

- On the Create Filesystem menu, clear the Create Automatic Snap FS Schedule (every 4 hours) checkbox.
- Modify the automatically created snapshot schedule to satisfy your requirements.

To avoid SAN filesystem inconsistencies, AxiomONE Storage Services Manager does not provide for creation of schedules for Clone LUNs (formerly called Snap LUNs). Before creating a Clone LUN for a LUN, be sure that the SAN host has placed the filesystems on that LUN in a state where they will be consistent. If desired, you can use the CLI with a scripting language to create Clone LUNs periodically.

9.26 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 an Administrative Action 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.

9.27 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.

9.28 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.

9.29 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 (filesystem or LUN) is assigned, the Pillar Axiom system recreates the volume at the new location. Attempting to use a logical volume while it is being moved can result in lost data.

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

9.30 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.

9.31 Running Slammer Diagnostic Tests from the GUI

When you run diagnostics on a Slammer through the GUI, read the instructions and warnings concerning the removal of external cables.

Note: Slammers always run diagnostics when powered on or restarted. Hence, you should not need to run diagnostics unless instructed to do so by a support engineer.

As the diagnostic executes, resources on the Slammer control unit (CU) will fail over to the other CU. The CU being tested will go offline, which generates an Administrator Action indicating that the CU has failed. The system status will show Critical.

If the diagnostic passes:

- The Slammer CU is placed online, the Administrator Actions are automatically deleted, and the system Status shows Normal.
- For SAN Slammer CUs, all resources will automatically fail back.
- For NAS Slammer CUs, if automatic failback has not been enabled, you need only execute the Administrator Action associated with the failed over CU to fail back the resources to the CU that has successfully completed diagnostics.

Important! Do not attempt to run Diagnostics on more than one Slammer CU at a time or when the system is busy. Doing so may cause additional resources to go offline.

9.32 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.

9.33 GUI Accurately Displays the Status of Components

In earlier releases, during system or core restart, all components would show a status of “Booting”. Beginning with release 2.0 of Pillar Axiom 500 systems (release 1.0 of Pillar Axiom 300 systems), the GUI more accurately displays the status of the components as they are discovered and initialized by the management software. For example:

- The typical initial display shows the active Pilot control unit (CU) as Booting. The standby Pilot CU may show Warning, Offline, or Booting, then transition to Online when start-up finishes.
- The initial Slammer state shows as Unknown and then proceeds to Boot State Ready, Booting, “Booting 0xnnn”, and then Online as the Slammers are discovered and initialized.
- The initial Brick state shows as “Offline” and then “Booting” as the storage enclosures are discovered and initialized.

9.34 A Disk Drive May Display Blank Data in the GUI

Disk drives are validated by the Pillar Axiom system. In some cases, the GUI may report a blank part number or serial number for a disk drive and, occasionally, a "Cannot read" status for that disk drive. However, the system will perform normally.

9.35 Changing Components

Use Guided Maintenance in the GUI when changing hardware components. Guided Maintenance provides instructions for you and performs tasks to get the system ready for the component replacement. Make sure to follow the instructions provided.

Notes:

- Adding memory to existing Slammer control units (CUs) in the field is not supported in this release. Contact your account representative for assistance.
- Replacing Brick and Slammer chassis are not supported in this release. If you have attempted to replace a Brick chassis, contact the Pillar World Wide Customer Support Center for assistance in recovering.

9.36 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 control unit (CU) and, depending on the activity on the Pilot 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.

9.37 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.

9.38 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.

9.39 Replacing a Disk Drive

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

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

Important! You should contact the Pillar World Wide Customer Support Center before pulling a disk drive.

- If a disk drive fails to be accepted into a Brick and the disk drive is set to Rejected status, do not attempt to use that disk drive. Contact Pillar Data Systems for another disk drive and for assistance.
- If an Administrator Action asking you to accept the disk drive is generated, be sure to select the Accept Drive option, which will initiate a copyback operation.

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

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

9.40 Moving Disk Drives

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

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

9.41 Reseat Disk Drives before Powering On New or Replacement Bricks

The disk drive latch may appear to be fully latched, but sometimes the disk drive is not making good contact with the Brick chassis midplane. With poor contact, the disk drive will fault, and the GUI will typically display a state of Unknown for that disk drive:

```
Drive 06  05_Fault      00_Unknown
```

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

If a disk 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 disk drive carrier will not be flush with the other carriers.

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

1. Press on the disk drive to latch them.
2. Press the disk 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 disk drive carrier unnecessarily. Doing so can lead to potential troubles in the future.

9.42 Testing Multiple Disk Drive Failures

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

9.43 ACT LEDs on Disk 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 disk 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.

9.44 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 filesystem or LUN is physically on the Brick.

9.45 Adding a Brick Generates Error and Warning Messages

When you add a Brick storage enclosure to an existing Pillar Axiom system, the system begins the process to bring the Brick online. While the system is bringing the Brick online, 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. During the bring-up process, the status of the Brick will go from red to yellow to green. After the system completes the process, the Brick will show a Normal status and will remove all Administrator Actions related to adding the Brick. If any Administrator Actions remain, contact the Pillar World Wide Customer Support Center.

9.46 Priority QoS Bands and Fibre Channel Bricks

As of 03.01.00, Fibre Channel (FC) Bricks support the bleed up of lower-performance bands from SATA to FC storage in a mixed SATA and FC system. Furthermore, you can create one logical volume having a non-Premium QoS that occupies all SATA and all FC Bricks. Logical volumes having a Premium QoS setting, however, are constrained to FC Bricks.

9.47 Testing RAID Rebuild and Simulated Drive Fault

You can use Guided Maintenance to identify a Brick and to show the location of an individual disk drive for testing disk 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 Administrator Action 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 “Replace Hardware” to remove the old drive and replace it from spares as instructed.

9.48 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 disk drives are functioning and Bricks remain operational.

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

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

When the disk drives are replaced, as soon as enough drives become available to allow the Pillar Axiom system to recreate the data LUN, the system will perform a `RecreateRAIDArray` task and create a new, blank data LUN to replace the faulted data LUN. All data on the original Brick LUN will be discarded.

Under some conditions, the Pillar World Wide Customer Support Center may be able to recover from the original faulted LUN. To do this, all internal fabric connections to the Brick must be disconnected during the recovery; otherwise the system will detect the recovered disk drives and proceed with the creation of a new blank LUN.

9.50 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 attempt the Verify function during Guided Maintenance of a Slammer motherboard. The verification will fail.

Important! Do not run Slammer diagnostics during motherboard replacement.

9.51 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.

9.52 Reverse Name Lookups for NFS Exports

NFS mount authentication time has been improved by adding reverse name lookups for NFS exports defined by hostnames. For this to be effective, you should configure external naming servers with the reverse records.

- DNS servers should be configured with PTR records for all client hostnames.
- NIS servers should be configured with host.byaddr records for all client hostnames.

Important: On the File Server Services tab, you should specify the appropriate order for “Host Name Resolution Search Order”. Name resolution policies must be applied consistently. Configure the reverse and the forward name lookup systems to provide matching names (both must return fully qualified names, such as “myhost.mycompany.com”, or both must return unqualified names, such as “myhost”). NFS mount requests may be rejected when a DNS server returns a fully qualified name and an NIS server returns an unqualified name.

For example, if the DNS server returns the host name as “myhost.mycompany.com”, the customer’s netgroup must have a membership list that includes the fully qualified name “myhost.mycompany.com”. Similar restrictions apply to name resolution by means of local files.

9.53 Deleting Files from a Full Filesystem with Snapshots

If a filesystem containing snapshots is allowed to completely consume the allocated space, it may become impossible to delete files from that filesystem to recover space.

Snapshots consume space from the original filesystem allocation in order to preserve QoS. If a file is deleted, the data from that file would need to be placed in the appropriate snapshot copies to preserve the integrity of existing filesystem snapshots. If there is no space left, the file deletion will be failed in order to preserve these views.

To delete files from such a filesystem, delete one or more snapshots to recover filesystem space, or allocate more space if there is storage available in the current or higher QoS pool.

9.54 Uploading Empty Files Not Allowed

You cannot upload empty files (zero bytes) to a Pillar Axiom system. This restriction applies to all system interfaces where a file upload is allowed.

9.55 Filesystem Checking (FSCK) and Consistency

In this release, FSCK is fully functional in repairing filesystems.

As an additional data protection measure, when a filesystem is created, the default behavior is to create a four-hour snapshot schedule for the filesystem. Pillar recommends these snapshots, because in rare circumstances, the FSCK may fail to repair the filesystem and will provide an Administrator Action choice to revert the filesystem to a snapshot or to delete the filesystem.

If the repair is unsuccessful and there are no available snapshots, you may still place the filesystem online for recovery of the data. Exporting such a filesystem read-only for this recovery is recommended. Contact the Pillar World Wide Customer Support Center if this occurs or before deleting a filesystem due to a failure of the filesystem check.

You have the option to place the filesystem online to revert to a known good snapshot to avoid a lengthy FSCK. Contact the Pillar World Wide Customer Support Center for assistance in determining candidates for good snapshots before reverting in this manner.

It is possible to recover from some filesystem conditions by placing the filesystem online and reverting to a snapshot. The recommendation is to export such a filesystem read-only for recovery of the data.

If the filesystem becomes inconsistent and a Snap FS exists, FSCK can revert the filesystem back to when the snapshot was created. Any data saved, stored, or modified on the filesystem after the snapshot was created will be lost.

The snapshot frequency defined by a Snap FS schedule is also important. The time lapse between when a Snap FS was created and the discovery of filesystem inconsistency determines how much data will be recovered should the filesystem need to be reverted.

If filesystem inconsistency is detected, the system performs the following:

1. Takes the filesystem offline so that no further changes can be made to the data.
2. Generates an Administrator Action that provides multiple options to the administrator:
    - Perform a filesystem consistency check.
    - Delete the filesystem.
    - Revert to a previous snapshot.
      If you choose this option, all changes to the filesystem between the time of that snapshot and the time the filesystem issue was detected is lost.

3. Checks the filesystem for consistency.
    - If the filesystem is consistent, the check is complete and FSCK automatically puts the filesystem online.
    - If the filesystem is inconsistent, FSCK takes one of these courses of action:
        - If the filesystem is fixable, FSCK fixes it and puts the filesystem online.
        - If the filesystem is unfixable, FSCK performs the following actions:
            - Reverts the filesystem to an earlier snapshot.
              In this case, FSCK reports “FSCK Complete” along with the snapshot ID.

              Note: Any data saved, stored, or modified between when the snapshot was taken and when the filesystem issue occurred will be lost.
            - Verifies the consistency of that snapshot.
              If it is fixable, FSCK fixes it and puts the filesystem online; otherwise, FSCK repeats the above steps.

    Note: If no good Snap FSs exist and the initial FSCK fails to recover the filesystem, Pillar recommends that you *not* delete the filesystem but instead contact the Pillar World Wide Customer Support Center. The Support Center may be able to help you recover the filesystem.

Once the consistency check is complete and the filesystem is online, end-users may need to remount the filesystem so they can reconnect.

We strongly recommend that you run only one FSCK process at a time. Even though multiple FSCK processes will queue, the best practice is to check a single filesystem and wait until that check completes before starting a subsequent FSCK process.

9.56 Replication of a SecureWORMfs Filesystem

When you target a SecureWORMfs filesystem for an AFR file replication operation, you must set the delivery mode under the file transfer options to “fast”; otherwise, AFR attempts to create a working file and rename it when the transfer is complete. Because SecureWORMfs filesystems are read only, that is not possible. Setting the file transfer options to “fast” prevents the creation and renaming of a working file.

9.57 Increasing Redundancy Requires More Space Than Requested

When increasing the Redundancy QoS parameter of a filesystem, the actual amount of additional space the system allocates for the increased redundancy is typically greater than what is requested. For example, you should expect that, when increasing the redundancy of a filesystem to Double, the system could allocate an additional 2 GB and, when increasing the redundancy to Triple, the system could allocate an additional 5 GB.

9.58 Creating Quotas on Non-Empty Directories

When creating a quota on a non-empty directory, you have a choice of allowing the filesystem to go offline temporarily or failing the quota request.

- *Allow the filesystem to go offline*. The system takes the filesystem offline temporarily to calculate quota usage of the directory. The system traverses the entire directory and counts the number of blocks consumed by the directory. The duration of the offline state of the filesystem depends on the number of objects contained in that directory.
- *On empty directories only*. The system tells you about any non-empty directories and fails the quota request in those cases. These quotas can be implemented when taking the filesystem offline is acceptable or by creating a new directory and quota and moving the data to the new directory.

9.59 CIFS Support

A File Server can act as a native member server to an Active Directory environment or emulate a Windows NT member file server.

In the Active Directory environment, Kerberos authentication is supported. A DNS server is required for name resolution, and a Domain Controller is required to provide Active Directory support and to act as the Kerberos Key Distribution Center (KDC). Customers that have implemented higher security policies should be aware that LDAP signing is not supported (the Pillar Axiom CIFS server cannot join the domain when ”LDAP server signing requirements = Require signing” is specified on the Domain Controller).

As a Windows NT member file server, NTLM authentication is supported. It uses NetBIOS naming services and requires a Domain Controller and WINS server with static IP addresses and legal NetBIOS names.

A File Server may be used in Windows 2000 or Windows 2003 domains with some special configuration requirements. These requirements are discussed in the *Windows Integration Guide for NAS Systems*.

Additional CIFS features will be supported in future releases.

9.60 *Kerberos Key Encryption in Windows 2008 Domains*

Default Kerberos encryption key length is 128 bits in Windows Server 2008. However, Pillar Axiom systems only support 64-bit keys. Therefore, when joining a Windows Server 2008 domain, you must set Kerberos so it uses 64-bit key lengths, which is equivalent to the NT security level.

9.61 Domain Local Groups Cannot Be Used in Mixed-Mode Domains

In Windows configurations that are using mixed-mode domains, if Access Control Lists (ACLs) of the domain local group are applied to share objects, access to the objects on the share can fail. This limitation is correct domain client behavior, as defined by Microsoft. See http://support.microsoft.com/kb/296369.

Here are some possible alternatives:

- Use domain global group ACLs instead of domain local ACLs on share objects.
- Refain from applying domain local ACLs on share objects.
- Upgrade the domain from mixed mode to native mode.

9.62 Create CIFS Share

This release implements the `\\<fileserver-name>\IPC$` administrative share, which allows CIFS clients to browse the filesystems that are shared by a File Server. Pillar Axiom systems do not automatically create the `C$`, `D$`, and `<share-name>$` style hidden administrative shares. If these shares are desired, share the filesystem as `C$`, `D$`, and `<share-name>$`. For example, to create a hidden administrative share for a filesystem named `foo`, create a share named `foo$`.

9.63 Sort Order on Pillar Axiom Systems Differs from Native Windows

When using `dir` to list directory contents, specify the requested sort order by means of the command parameters. For example, to order by name, use:

```
dir /on
```

Some CIFS clients, such as Windows Explorer, sort lists of files and directories in a specific order. Other applications, like the `dir` command, default to return the list in the internal order maintained within the File System.

The Windows NTFS File System appears to maintain files sorted alphabetically. The Windows FAT File System does not. The Pillar Axiom File System also does not maintain a default alphabetical sort sequence.

As a result, the Pillar Axiom sort order may differ from native Windows when viewed through some applications. A list of files produced with the default options of the `dir` command do not return a list sorted alphabetically.

9.64 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 pdscli, 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.

9.65 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.

9.66 SMI-S Should Not Be Enabled on 512 MB Pilots

Important! By default, all 512 MB (older, non-RoHS) Pilots have SMI-S disabled on startup. Because of limited memory, you should not enable SMI on these Pilots. If SMI-S is disabled and you would like to run it, contact the Pillar World Wide Customer Support Center for assistance.

9.67 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.

9.68 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; if the failures occur during startup, the system marks the CU as Excluded. This behavior may be triggered by repeatedly testing power failure simulation for Slammers.

If the system completes startup successfully, it will attempt to remove each Excluded Slammer control unit (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 CU is part of a NAS Slammer and automatic failback is *not* enabled, the system attempts to transition the CU to Failed Over and leave it there for recovery.
- If automatic failback of NAS Slammer resources *is* enabled, the system attempts to transition the CU from Excluded to Failed Over to Failback and, if that succeeds, the system puts the CU online.
- 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.

9.69 Hardware Lockout Due To Repeated Power Cycling

IMPORTANT! Do not repeatedly power cycle a Pillar Axiom system or any of its hardware components. Doing so may automatically trigger the hardware-fault lockout mechanism. 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.

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

9.70 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 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.

9.71 Battery Removal

When removing a Slammer battery, be sure to use Guided Maintenance. After you click the Prepare System button in the GUI, Guided Maintenance prepares the system for replacement of the battery:

- Flushes cached data to the Bricks.
- Places the target control unit (CU) 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.

9.72 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.

9.73 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.

9.74 Preventing Filesystem Corruption Due To Multiple Component Failures

In the unlikely event that more than one disk drive has failed, filesystem corruption may occur. To avoid this possibility, take advantage of the redundancy and availability options available when creating filesystems.

9.75 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.

9.76 Feature License Reconfiguration

When you change the feature license configuration on a Pillar Axiom system, the AxiomONE Storage Services Manager (GUI) will restart and you will need to log back in. Wait for several minutes before logging in to give the system time to reconfigure; otherwise, you may receive a Page Not Found error.

9.77 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 sec. If the AxiomONE Storage Services Manager does not show updated states correctly, refresh your browser display.

9.78 NDMP Backup Operations on Full Filesystems

As an NDMP backup operation begins, the system takes an immediate snapshot of the filesystem. This snapshot is the NDMP data source for the backup. At the end of the backup operation, this snapshot is deleted. Working from this snapshot allows clients to use the filesystem during the NDMP backup operation.

If the filesystem is full, the backup operation may fail or it may be impossible to delete files from the original filesystem during the backup operation, because the data from those files would need to be stored in the backup image snapshot.

The space allocated to the filesystem should be increased to allow any data that is modified or deleted during backup operations to be safely stored in the snapshot created for the backup. If it becomes necessary to recover space, delete any unnecessary snapshots and wait for the space to become available. If there are no snapshots, delete at least 16 KB of data prior to running the backup.

9.79 NDMP DMAs May Disable Tape Storage Devices During a Software Update

Updating system software may cause tape drives or libraries attached to a Pillar Axiom system to go offline.For example, VERITAS has designed their products to take the tape storage devices down the first time a problem appears. In this case, use the administrative tools provided by your NDMP-based data management application (DMA) to bring the tape storage devices back online. If you need assistance, contact the Pillar World Wide Customer Support Center.

9.80 Time Zone Must Be Reset after a Software Upgrade

As of release 02.06.00, the time zone fields were updated to provide a more comprehensive and descriptive list of time zones. Pillar recommends that, after the upgrade completes, you reset the time zone for your system through the Modify Time Settings action in the System > Summary content pane. To reset the time zone, select the appropriate entry in the Time Zone field. This action will update the saved time zone values to the new field choices.

Resetting the time zone need be done only once when upgrading your software to release 02.06.00 or higher.

9.81 Point-in-Time Volume Copy Space Allocation

The system can use point-in-time Volume Copies to support NDMP backup requests against non-WORM filesystems. When using an NDMP Volume Copy, there is an additional space limitation above the amount required for file-based NDMP backups. The Volume Copy implementation requires an additional amount of storage equal to the maximum size of the filesystem being backed up.

9.82 Jumbo Frames

To implement jumbo frames, be sure that all Ethernet NICs in all systems, Ethernet switches, local router interfaces, and all other network devices on the same local networks as the Pillar Axiom system are capable of supporting jumbo frames and are configured for the same effective MTU.

Refer to your switch documentation for information on prerequisite software and firmware and for instructions on configuring the switch for jumbo frames. Refer to the documentation for the NIC interface in all client systems and other network devices for information and restrictions on configuring jumbo frames.

The performance boost with jumbo frames will be most noticeable for client systems with slower processors or interrupt handlers that may benefit from the lower interrupt rate offered by jumbo frames. The increase in performance will be most noticeable for single-client stream data transfers.

9.83 Link Aggregation over Fast and Gigabit Ethernet Switches

Link aggregation groups several network links into a single logical link. Link aggregation provides load balancing and fault tolerance for multiple Ethernet links. Using extensions of Ethernet auto-negotiation, the link partners exchange MAC control frames to provide for load balancing and link fault detection.

To make use of link aggregation on a Pillar Axiom system, be sure that:

- All ports that will be used for a given logical aggregated link are on the same Slammer control unit (CU). Ports from multiple Slammers or Slammer CUs must not be configured into the same aggregation group.
- Many Ethernet switches require that all links in an aggregated link be on the same blade or Ethernet controller chip. Consult the switch vendor's documentation for restrictions.
- The Slammer connects to a 100/1000 BaseT switch that has Auto-speed enabled.
- The Ethernet switch must conform to the IEEE 803.2ad link aggregation standard and use link aggregation control protocol (LACP) for managing the links.
- The Ethernet switch should be configured to actively advertise link aggregation capability. The Slammer ports will respond to but not advertise link aggregation.

Note: To provide for continued operation in the event of a Slammer CU failover, if link aggregation is enabled on one Slammer CU, it should also be configured on the partner Slammer CU.

9.84 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.

9.85 OpenSSH Has Some Vulnerabilities

The Pilot management controller uses a version of OpenSSH (OpenBSD Secure Shell) that has several known vulnerabilities. However, for someone to gain access to a Pilot, they would need to successfully perform the following actions:

1. Gain access to the data center network.
2. Install an SSH Encrypted License from Pillar Data Systems using the Primary or an Administrator 1 account.

3. Enable SSH access for the Support account using the Primary or an Administrator 1 account.
4. Use the Support Administrator login credentials to gain access to the Pilot management controller directly by means of SSH.

It is unlikely that anyone would be successful in performing those four actions. Furthermore, 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.

9.86 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. Before attempting a shutdown or restart, ensure that no background tasks are running. If you are unable to cancel a task, contact the Pillar World Wide Customer Support Center.

9.87 Login to Oracle EM Plugin Fails

The login to the Oracle EM plugin 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.

9.88 Changes to Host Associations and LUN Mappings

Release 02.07.00 introduces many changes to host associations and LUN mappings for the AssociateInitiatorsToHost feature and to APM.

The GUI page for PeformAssociateInitiatorsToHost enforces many of the changes. For example, Modify can no longer be used to change the name of an Associated host. It also prevents the user from associating an already associated Initiator/port without first removing its association.

9.88.1 Definitions

Initiator. Equivalent to a port (WWN) when using FC or IQN when using iSCSI.

Wrapper host. A host with a single Initiator and has the same name as the Initiator. It is created internally by the Pillar Axiom system when an Initiator is discovered.

Associated host. A host that the user has created with the PerformAssociateInitiatorsToHost request (one or more Initiators). This request is referred to as *PAITH*.

APM host. A host created from a PerformConfigureSANHost request from APM (one or more Initiators). This request is referred to as *PCSH*.

The Pillar Axiom system always creates a Wrapper host for all discovered Initiators and moves each according to Associations and/or APM relevance. A user can map to all of these hosts.

9.88.2 Behavior

9.88.2.1 Moving an Initiator from a Wrapper Host

When an Initiator belonging to a Wrapper host is moved into another host by either a PAITH or PCSH request, the LUN mappings of the Wrapper host are moved to the new host and removed from the Wrapper host. The Wrapper host is deleted. The original mappings belonging to the new host are unaffected. Mappings from the Wrapper host that conflict in LUID or Number with mappings in the new host are deleted and a LUNMappingDeleted event is generated.

9.88.2.2 Moving an Initiator from an Associated Host

When an Initiator is moved from an Associated host by a PAITH request, the mappings of the previous owner Associated host are removed from the Initiator, and the Initiator is mapped with the mappings of the new owner Associated host. The old owner Associated host (even if it has no Initiators left) is left with its LUN mappings in place. No mappings associated with any old owner Associated host are moved to the new Associated host.

The old host must be manually deleted if it's no longer needed.

9.88.2.3 Using PCSH to Move an Initiator to an APM Host

An Initiator can be moved to an APM host from an Associated host or from another APM host by a PCSH(V2) request. For the previous Associated host, its mappings are moved to the new APM host. For the previous APM host, its mappings are not moved to the new APM host.

Important! When moving an initiator to an APM host, APM on the new host (and on the previous owning APM host, if applicable) must be able to log in to the Pillar Axiom system and achieve a "Communicating" status in the AxiomONE Storage Services Manager (see Storage > Hosts display screen). In this way, APM can exchange the necessary information with the Pillar Axiom system:

- The new APM host can report the initiator to the system (but not get the LUNs unless they too have been moved).
- If applicable, the previous owning APM host can remove the initiator(s) when it logs in to the system.

9.88.2.3.1 Moving an Initiator from an Associated Host to an APM Host

Mappings from the Associated host that conflict in LUID or Number with mappings in the new APM host are deleted and a LUNMappingDeleted event is generated. A LUNMappingCreated event is generated for each mapping successfully moved. All of the LUN mappings of the previous owning Associated host stay in place for that host, even if no Initiators are associated with the host.

The old Associated host stays in place until the user deletes it.

9.88.2.3.2 Moving an Initiator from an APM Host to another APM Host

The old mappings to the Initiator are removed and the Initiator is mapped to the configuration of the new owning APM host.

9.88.2.4 Combinations

For combinations, the processing goes from Wrapper host to Associated host to APM host with no order within the different types.

Example:

If an APM claims ownership of Initiators from a Wrapper, Associated, and APM host at the same time, mappings are moved from all Wrapper hosts to the APM host first, followed by moving mappings from the Associate hosts. Then the APM Initiator is moved and remapped.

9.88.2.5 Notes

9.88.2.5.1 Deleting a Host

When deleting a Wrapper, Associated, or APM host, all Initiators that are connected are wrapped with a Wrapper host and all of the LUN mappings belonging to the host are preserved with the new Wrapper host. No Wrapper host is created for a disconnected Initiator.

9.88.2.5.2 Renaming an Associated Host

To rename an Associated host while preserving the mappings, the Associated host must be deleted and then the Initiators can be associated with the new named host.

9.88.2.5.3 Renaming an APM Host

To rename an APM host while preserving the mappings, the APM driver must be stopped, the host deleted, and then the APM driver started with the new name.

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 300, 500, and 600 system software.

10.1 *Pillar Axiom Administrator's Guide*

- In Chapter 2, "Free Capacity and Volume Creation", page 46, for the paragraph that begins with "As a general rule…," add the following to the end of the sentence:

  "…, minus system overhead. This system overhead is as follows:

  - *"50 GB for a system containing only SATA Bricks or only Fibre Channel Bricks.*
  - *"100 GB for a system containing a mix of Brick types."*

- In Appendix A, page 246, add the following information under Description:

  "SecureWORMfs objects are not intended to be dynamic production filesystems. Instead, WORM filesystems are intended to be used for archival purposes. As such, when you enable a WORM filesystem to protect files when the protocol closes them, the filesystem protects those files immediately, not allowing any further modifications. As part of fulfilling a user request to open a file, some applications open the base file, copy the contents to a temporary location, and then close the base file while allowing the user to work on the contents of the temporary file. In this case, because the WORM filesystem immediately protects the base file, which was just closed, an attempt to save the edits to the base file will fail.

  *"**Tip:** You can enable the feature of protecting a file immediately when it is closed by creating a directory called .pillar_doPoC (case insensitive) in the root of the WORM filesystem. You can disable this "protect on close" feature anytime by removing this directory.*

  *"**Note:** WORM filesystems that do not have file protection enabled in this way will not protect those files immediately when they are closed but will protect them within 24 hours."*

- In Appendix A, page 246, replace the second sentence in the first paragraph that describes the Protected File Integrity Scan field:

  "This scan occurs every 90 days beginning on the day in which the SecureWORMfs was created and starts at the time set in the Daily File Auto-Protection and Compliance Scan."

10.2 *Pillar Axiom 600 Service Guide*

On page 269, add the following notice:

> *"When you intend to use non-Pillar PDUs, check first with your Pillar account representative to ensure you do not jeopardize your system warranty by installing the non-Pillar PDUs."*

10.3 *Pillar Axiom 600 Advanced Hardware Installation Guide*

In Chapter 2 on page 22, expand the **Important!** notice by integrating the following information:

> *"We recommend that Pillar racks be used to install Pillar Axiom hardware components. When using non-Pillar racks, do not use Telco two-post racks. Instead, use a four-post rack that can support the weight load of a Pillar Axiom system. Additionally, be sure the non-Pillar rack has square mounting holes in the vertical channels. Round mounting holes are not acceptable."*

Modify the row for "Vertical Channels" in Table 50 on page 175 as follows:

> In the "Pillar rack" column, change the first bullet to "Square hole unthreaded".
>
> In the "Non-Pillar rack" column, change the last bullet to "Square EIA-standard mounting holes required".

We recommend that Pillar racks be used to install Pillar Axiom hardware components.

10.4 *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`.

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

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

- *AxiomONE Path Mananger 3.0 Installation Guide and Release Notes for RHEL 5.2*
- *AxiomONE Path Mananger 3.0 Installation Guide and Release Notes for OEL 5.2*
- *AxiomONE Path Mananger 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.

10.6 *SAN Attachment SCSI Command Reference*

In Table 76 "Target Port Group Descriptor" (page 108), bit 7 in byte 0 is now the PREF bit. Change the first row to the following:

Byte	Bit							
	7	6	5	4	3	2	1	0
0	PERF	Reserved			Asymmetric Access State			