Pillar Axiom® 300, 500, 600

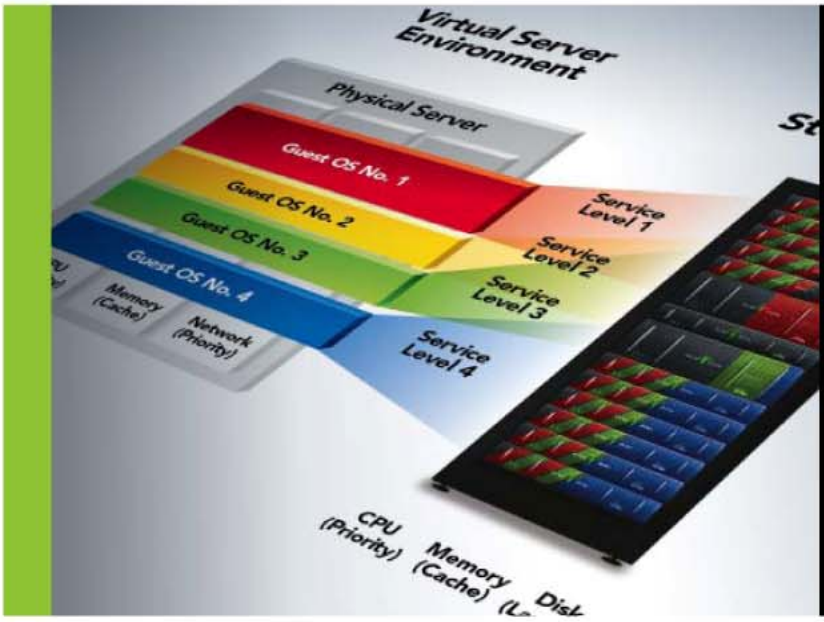

Customer Release Notes

For Release 4.3.1

pillar DATA SYSTEMS

Pillar Axiom 300 **Pillar Axiom 500** **Pillar Axiom 600**	**Document Number: 4420-00026-3500 Rev 1.1** **Document Title: Customer Release Notes**

Revision History

Rev Description	Rev Date	Effective Date
Release 4.3.1 GA Rev 1.1	2011-02-08	2011-02-11
Release 4.3.1 GA	2011-01-03	2011-01-19
Release 4.3.0 Beta	2010-12-14	2010-12-28
Release 4.2	2010-06-21	2010-07-22
Release 4.1	2010-03-30	2010-04-27
Release 4.0	2009-11-20	2009-12-21
Release 3.4 GA	2009-07-22	2009-08-28 [Pillar Axiom 600 only]
Release 3.3 GA	2009-03-20	2009-05-01
Release 3.2 GA	2008-10-20	2008-10-27
Release 3.2 Beta	2008-10-10	2008-10-20
Release 03.01.00	2008-06-06	2008-06-18
Release 03.00.00	2008-03-12	2008-03-14
Release 02.08.00	2008-01-10	2008-01-14

1 Terms and Conditions of Use

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

2 Purpose

This document describes new features, capacities, configuration requirements, operating constraints, known issues and their workarounds, and other items for release 4.3 of all Pillar Axiom Storage System platforms. 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.

3 Product Release Information [a]

Release 4.3 is a maintenance and feature enhancement release of the unified NAS/SAN Pillar Axiom software for all Pillar Axiom platforms.

3.1 System Enhancements

This update provides substantial quality improvements to all NAS and SAN Pillar Axiom models along with several feature and performance enhancements.

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 4.3.

For the list of defects that this release resolves, see Table 11 beginning on page 37.

3.1.1 Hardware Enhancements

- Version 2 Fibre Channel (FC) Bricks are now available that support 4 Gb/s optical connections to the Slammers in addition to legacy 2 Gb/s copper connections. The 4 Gb/s optical capability doubles the backend bandwidth that 2 Gb/s copper connections provide.

  Note: All Version 2 FC Bricks have RAID controllers and do not offer the Expansion Bricks option that was available for Version 1 Bricks.

- 10 GbE network interface modules (NIMs) are now available for Pillar Axiom 600 Series 3 Slammers. These NIMs (copper and optical) provide high-speed, front end connectivity to NAS data networks. As in previous GbE NIMs, this 10 GbE version provides a Fibre Channel (FC) HBA option for connecting to tape devices in support of NDMP backups.

  Important! Consult the vendor's requirements for cabling in your 10 GbE environment. For information regarding supported FC and copper cabling, refer to Section 9.82 in this document.

- 8 Gb/s FC NIMs now support 1 Gb/s iSCSI connections.

3.1.2 Software Enhancements

- Backup administrators can now indicate the preferred Slammer control unit (CU) that the system should use to perform an NDMP backup. By configuring the data management application (DMA) so it uses a specific TCP port, a backup administrator indicates the preferred CU through which the backup should be performed. Port specification permits the backup administrator to ensure that NDMP backups run on the same Slammer CU on which the filesystem being backed up is homed.

  An administrator determines this port through use of the `pdscli GetAllSlammers` command. This command returns an NDMPPortID value for each Slammer control unit (CU) configured in the system. This value is the TCP port number associated with that CU. For example, if `GetAllSlammers` returns TCP port 10003 for Slammer2 CU0, you can specify port 10003 in the DMA when backing up a filesystem homed on that CU.

[a] Throughout this document, all references to release 2.x apply to the Pillar Axiom 500 system. The corresponding release number for Pillar Axiom 300 systems is 1.x. In other words, the software for those two release numbering schemes is equivalent.

- New CIFS security features, which can be enabled through the CLI or the GUI by configuring the File Server, are now available:[b]
    - **Access based enumeration (ABE)**, when enabled, improves security by not displaying filesystem objects for which the client has no access rights. An ACL controls what a client can do, whereas ABE controls what the client can see.
    - **Bypass traverse checking (BTC)**, when enabled, emulates the behavior of a standard Windows server, allowing users to traverse to directories without requiring that they have rights to view all directories in the path. Access rights, however, are still checked for the target object itself.
- The maximum length of the names for filesystems, Snap FSs, and Clone FSs has been increased to 63 bytes (or 21, 3-byte UTF-8 characters).

Tip: For optimal performance with the 4.3 release of the Pillar Axiom software, Pillar recommends that APM clients use round-robin access from the host. Doing so ensures that all ports and Pillar Axiom CPUs are fully utilized.

3.2 Changes to How the Pillar Axiom System Operates

3.2.1 Capacity Utilization Across Brick Types (Storage Classes)

In software releases prior to release 4.1, for a Pillar Axiom system containing a mix of SATA and Fibre Channel (FC) Bricks, individual logical volumes could utilize space from both the SATA and FC storage pools. As of release 4.1 and later releases, administrators cannot *create* volumes that span Brick types (called *Storage Classes* as of release 4.1). Legacy volumes created prior to release 4.1 that span Storage Classes, however, can continue to exist in the system after updating to release 4.3.

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

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

The system-wide available capacity, however, continues to be reported in the same manner as it was reported in earlier releases.

3.2.2 Get Storage Configuration CLI Command

Starting with release 4.1, the CLI uses a new command `GetStorageConfig` rather than `GetStorageConfigDetails` (which is no longer supported) to query for the system capacities. For `GetStorageConfig`, the total system capacity is reported just as before. However, instead of reporting capacity on the basis of performance bands, the system now reports capacity on the basis of Storage Classes.

[b] Instructions to enable the ABE and BTC features are provided in Section 9.66.

3.2.3 SecureWORMfs Daily Scan

In software releases prior to release 3.3.26, the Daily Protection Scan feature scanned all protected files on the WORM filesystem on a daily basis. Starting with release 3.3.33, however, the Daily Protection Scan feature was removed. The following mechanisms take the place of the Daily Protection Scan feature:

- For CIFS connections, the system automatically protects files in this way:
    - Files less 10 MB in size are protected immediately, in-line.
    - Files 10 MB in size or larger are protected near in-line, but without holding up the client.
- For NFS connections, you can edit the Immutable field in the appropriate XML file in the `.attributes` directory that is located within the WORM filesystem.

3.2.4 Triple Redundancy

Beginning with release 3.0, the Pillar Axiom system dropped the Triple Redundancy Quality of Service (QoS) option for *new* volumes. Using double redundancy in a Pooled RAID 10 array, however, effectively provides quadruple redundancy, if that is desired.

See Section 3.4.1.1 for information about migrating logical volumes from triply redundant to doubly redundant.

3.2.5 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) or in the Command Line Interface (CLI). As such, administrators need to be aware of the following when updating a pre-3.0 Pillar Axiom system to release 4.3:

- 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 allow for the creation of inactive clones.

As of release 4.1, the functional equivalent of volume backups became inactive Clone FSs and inactive Clone LUNs (formerly referred to as *Snap LUNs*). These inactive objects have the same structure and behavior as their active clone counterparts, the difference being that inactive clones cannot be made accessible for I/O, while active clones are accessible for I/O.

The main purpose of an inactive clone, like that of its predecessor the volume backup, is archival. In the GUI, an administrator can create clones of either an active or inactive type. The CLI requests for clones (these requests for LUNs still use the term *snap* instead of *clone*) generally can be used to create or manipulate either active or inactive clones.

In release 4.1, as a result of Volume Backup and Copy no longer being supported, the associated Command Line Interface (CLI) requests were removed from the CLI. (Refer to the current CLI documentation to determine which requests are supported in this release.)

See Section 9.20 for a complete list of the commands that have been removed.

3.2.6 Maximum MTU for File Servers

In release 4.3, the maximum frame size or MTU (maximum transmission unit) property of a File Server has been reduced from 16,362 to 9216 bytes.

3.3 Available Licenses

Table 1 Feature licenses

License	Comments	Pillar Axiom configurations			
		300	500	600	600e
1 MB Stripe	Automatically included. Enables use of the Automatic Storage Management (Oracle) performance profile.	•	•	•	•
NAS Asynchronous Remote Replication	Optional: – This license must be installed on at least one of the Pillar Axiom servers hosting a replication pair. – When this license is installed, Professional Services must be engaged to establish a new installation or to migrate from an existing replication solution.	•	•	•	•
SAN Asynchronous Remote Replication	Optional. – This license must be installed on the Primary Pillar Axiom server hosting the source volume. – When this license is installed, Professional Services must be engaged to establish a new installation or to migrate from an existing replication solution.	NA	•	•	•
AxiomONE Path Manager	SAN Slammer: Automatically included.	•	•	•	•
CIFS Protocol **NFS Protocol**	NAS Slammer: One protocol included at no cost.	•	NA	NA	•
	NAS Slammer: At least one protocol license must be ordered.	NA	•	•	NA
Clone FS (Partial image technology; read and write access)	Optional. For Pillar Axiom 300 and Pillar Axiom 600e, see “Copy Services bundle”.	NA	•	•	NA
Clone LUN (Partial image technology; read and write access)	Optional. For Pillar Axiom 300 and Pillar Axiom 600e, see “Copy Services bundle” below.	NA	•	•	NA

License	Comments	Pillar Axiom configurations			
		300	500	600	600e
Copy Services bundle	Optional: – NAS: Clone FS + Volume Copy / Backup – SAN: Clone LUN + Volume Copy / Backup	•	NA	•	•
Fibre Channel Protocol **iSCSI Protocol**	SAN Slammer: One protocol included at no cost.	•	•	•	•
NDMP	NAS Slammer: – Automatically included for Pillar Axiom 500 and Pillar Axiom 600 systems. – Must be purchased for Pillar Axiom 300 systems.	•	•	•	•
Pooled RAID 10	Automatically included.	•	•	•	•
Restricted System Brick	Restricts the Brick configuration to two.	NA	NA	• [c]	•
Extended System Brick [d]	Allows a Brick configuration of up to 64.	NA	NA	•	NA
SecureWORMfs	Optional.	•	•	•	NA
Snap FS (Partial image technology; read only access)	NAS Slammer: Automatically included.	•	•	•	•
SNMP	Automatically included.	•	•	•	•
Support Tool	Automatically included.	•	•	•	•
Thin Provisioning: FS	Optional.	•	•	•	•
Thin Provisioning: LUN	Optional.	•	•	•	•

[c] A Pillar Axiom 600e that is upgraded to a Pillar Axiom 600 configuration has the Restricted Brick license in addition to the Extended Brick license, both of which will be displayed in the GUI.

[d] When a Pillar Axiom 600e is upgraded to a Pillar Axiom 600 configuration, both the Restricted and Extended System Brick licenses can co-exist. In this case, the Extended System Brick license supersedes the limitations of the Restricted System Brick license. The Extended System Brick license is available only on Pillar Axiom 600 configurations that ship from the factory with release 4.2 or higher, or are upgraded from 600e configurations.

License	Comments	Pillar Axiom configurations			
		300	500	600	600e
Volume Copy / Backup (Full copy technology; read/write access)	Optional.	NA	•	•	•

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 system.[e] 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.

3.4 Pillar Axiom Software Update

The 4.3.1 release may be installed through the GUI Software Update process.

Note: If you have a Pillar Axiom system running software below release 2.4.2 and it contains Fibre Channel Bricks, contact the Pillar World Wide Customer Support Center to update the system.

Tip: If you have existing NAS systems running software earlier than release 3.0, contact your Account Representative to obtain a Thin Provisioning: FS license. Without this license, existing filesystems will continue to be accessible after the software update, and growth will continue to take place, *but newly created* filesystems will not be allowed to grow.

When updating a Pillar Axiom system from release 4.0 to release 4.3, *the update* is nondisruptive.

Important! Updating a Pillar Axiom system to release 4.3 from a release earlier than release 4.0 *is* disruptive to the data path. Before beginning the update process to release 4.3 from releases earlier than release 4.0, thoroughly read all of Section 3.4.

Important! For NAS systems running AxiomONE NAS Replication, both systems must be running on release 4.2.18 (or higher); otherwise, the target filesystem can become unusable. To ensure the target filesystem remains usable, follow the steps below.

1. Stop all NAS replication processes that are running.
2. Update all Pillar Axiom systems participating in NAS replication to release 4.3, as described in Sections 3.4.1 and 3.4.2.
3. Resume the replication processes.

Important! All customers who purchase NAS or SAN Asynchronous Remote Replication licenses must contact Professional Services to establish a new installation or migrate from an existing replication solution.

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

Important! After a NAS system is updated to release 4.3, that system should only be downgraded to release 4.2.18 or higher. Downgrading from release 4.3 to a 4.x release earlier than 4.2.18 will result in File System Inconsistency. Before attempting to downgrade, please contact the Pillar World Wide Customer Support Center for assistance.

If you prefer, you can have Pillar Professional Services update your system software. For more information, refer to Table 5 Contact information.

3.4.1 Prerequisites to Updating Pillar Axiom Software

3.4.1.1 Migrating Volumes From Triple Redundancy to Double Redundancy

Beginning with release 4.0, the Pillar Axiom system does not allow triply-redundant logical volumes. Those volumes must first be migrated to double redundancy in release 3.3 (or higher) before upgrading to release 4.3.

Even after a triply redundant logical volume that has clone storage has been migrated to double redundancy, the volume cannot be upgraded to release 4.3. Migrating such a volume to double redundancy only lowers the redundancy of the volume. To remove the triply redundant clone storage, delete all clones of the volume, and then set the clone storage capacity of the volume to zero, which causes the clone storage to be deleted.

Contact the Pillar World Wide Customer Support Center for assistance in migrating triply redundant volumes to double redundancy.

3.4.1.2 Minimum Required Release Level for Updating

Release 4.0 no longer supports triple redundancy volumes, including the internal volume used for the system configuration. An upgrade to any release 3.3 version produces an Administrator Action for any user volumes configured with triple redundancy and attempts to convert the internal system configuration volume from double to triple redundancy. The Pillar Axiom system cannot perform this migration if there is insufficient free space.

A pre-upgrade system configuration audit of all release 3 systems is required to ensure that there are no remaining triple redundancy volumes or other system configuration issues that could prevent a successful upgrade. Please contact the Pillar World Wide Customer Support Center for details.

Upgrades from lower releases to release 4.0 and higher must be performed disruptively. Release 03.03.25 or higher is required to ensure that the upgrade to release 4.x is done properly. All systems must be upgraded to 03.03.25 or higher before they can be upgraded to release 4.x. Refer to Section 2.4 (in particular Section 2.4.2.2) in the Release 3.3 Customer Release Notes (part number 4420-00026-2700) for more information.

If your system is running software release 4.0.0 or 4.0.1, you must first update your system to a minimum level of 4.0.2, which is non-disruptive to the data path.

3.4.1.3 Maximum Export Path Size

Beginning with release 4.0, the maximum sizes (in UTF-8 bytes) of export paths and the host names in the export access lists has changed:

Table 2 Changes to export path and host name limits

Object	Previous maximum (UTF-8 bytes)	Current maximum (UTF-8 bytes)
Export path	3072	1023
Host name	768	255

Important! We strongly recommend that, before you upgrade your system to release 4.3 from release 3.3.25 (or higher 3.3.x version), you ensure that all export and host names do not exceed these new limits. All exports or hosts with names that exceed these limits will become inaccessible after the update.

3.4.1.4 Maximum Length of Filesystem Names

As of release 4.3, the maximum length of filesystem names has increased to 63 bytes. When using 3-byte UTF-8 characters, the limit is 21 characters. This maximum name length applies to Snap FSs and Clone FSs as well.

3.4.1.5 Additional Intra-Control Unit Cabling for Slammers

Tip: For *single*-Slammer Pillar Axiom 600 systems, we highly recommend 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 4.3, you will receive an Administrator Action reminding you to provide that cabling.

3.4.2 The Release 4.3 Update Process

To update your Pillar Axiom system to release 4.3, you must first perform several additional preliminary steps. When those steps are complete, you can begin the update process.

3.4.2.1 Before Starting the Update

Before starting the software update process to release 4.3, perform the following tasks as needed:

- 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.
- If the AxiomONE Replication for SAN utility has been installed, before you upgrade the software on any of your Pillar Axiom systems, be sure to isolate all active replication pairs that might be using those systems as their primary or secondary system. If you upgrade without isolating pairs, replication will be disrupted and you will need to recreate any pairs that were left running.

For best results, isolate all pairs before you perform the upgrade. Because some restoring checkpoints may still be in progress, wait until all checkpoints have reached the history state before you begin the upgrade. Finally, after the upgrade, restart the pairs once operation is resumed on the upgraded systems.

- If the AxiomONE Replication for NAS utility has been installed, before you upgrade the software on any of your Pillar Axiom systems, be sure to follow the instructions outlined on page 8 for those systems running NAS replication.

  CAUTION! All systems participating in NAS replication must be at 4.2.18 (or higher); otherwise, the target filesystem can become unusable.

When updating a system from 3.x to 4.3, contact the Pillar World Wide Customer Support Center for instructions on how to perform a system audit, which is necessary when updating from release 3.x.It is not necessary to disconnect channel attachments from the Pillar Axiom system.

3.4.2.2 Starting the Update

When you start the update process from release 3.3.25 (or higher) to release 4.3, the Pillar Axiom automatically selects all the check boxes, including the Restart System option, if required.

Note: The Pillar Axiom system will automatically select the check boxes for upgrade and instruct the administrator not to change any of the selections without first contacting the Pillar World Wide Customer Support Center for assistance. This includes the Software Components and the Restart System option check boxes.

Important! The Restart System check box is available when you are performing a non-disruptive update from release 4.0.2 or higher to 4.3. However, do not select it.

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

3.4.2.3 After the Update

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

Important! After updating the system software to release 4.3, you cannot reverse the update and downgrade to an earlier release level without explicit action by the Pillar World Wide Customer Support Center.

3.4.3 Using the AxiomONE Replication Feature

If you have purchased licensing for AxiomONE Replication for NAS or SAN, or both, you must use Professional Services to establish a new installation (or migrate from an existing replication solution).

3.4.4 Software Versions for This Release

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

Table 3 Pillar Axiom Software versions

Software module		Pillar Axiom model	Software version
Pilot OS		All	04.03.00
Pilot Software		All	04.03.01
Slammer PROM		300 and 500	03.00.00
		600	04.00.00
Slammer Software		All	04.03.01
Serial ATA (SATA) Brick Firmware	SATA RAID controller, Version 1	All	07.19.19
	SATA V2 RAID controller, Version 2 (includes SSD)	All	00.19.19
Fibre Channel (FC) Brick Firmware	FC RAID controller, Version 1	All	00.79.19
	FC V2 RAID controller, Version 2		02.19.19
Brick Disk Drive Firmware		*Version depends on drive capacity and whether it is the spare or an array drive.*	

Note: If your Pillar Axiom system does not have Fibre Channel, SATA, or SSD Bricks installed, the GUI will not display the firmware entry for those Bricks.

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 4.3, be prepared for a change in how WWNs are managed. For more information, see Section 9.19 on page 52.

3.5 AxiomONE Path Manager Software

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

- 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 4.

Table 4 APM support

Operating System	OS Release	APM Release	Platforms
AIX	5.2	2.1	All platforms
	5.3 6.1	3.0	All platforms
Citrix XenServer	5.6	3.0	All platforms
Community Enterprise Operating System	4.5 and 4.6	3.0	32-bit x86, 64-bit x86, Itanium
	4.7	3.1	32-bit x86, 64-bit x86

(CentOS)	4.8	3.2	32-bit x86, 64-bit x86
	5.2	3.0	32-bit x86, 64-bit x86
	5.3	3.1	32-bit x86, 64-bit x86
	5.4	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86
HP-UX	11i v1 11i 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 Linux (OL) *Previously known as* 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
	4 8	3.2	32-bit x86, 64-bit x86
	5.2	3.0	32-bit x86, 64-bit x86
	5.3	3.1	32-bit x86, 64-bit x86
	5.4	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86
Oracle VM Server for x86	2.2	3.1	All platforms
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
	4.8	3.2	32-bit x86, 64-bit x86
	5.0	2.0	32-bit x86, 64-bit x86, Itanium

	5.2	3.0	32-bit x86, 64-bit x86
	5.3	3.1	32-bit x86, 64-bit x86
	5.4	3.2	32-bit x86, 64-bit x86
	5.5	3.3	32-bit x86, 64-bit x86, POWER
Solaris	8	01.04.05	SPARC with non-Sun HBAs, Axiom 500 only
	9	01.04.05	SPARC with non-Sun HBAs, Axiom 500 only
	9	2.0	SPARC with Sun HBAs
	10	2.0	SPARC, 64-bit x86
Windows	2000 Server	2.3	All platforms
	Server 2003 Server 2003 R2 Server 2008 Server 2008 R2	3.3	All platforms

Tip: For optimal performance with the 4.3 release of the Pillar Axiom software, we recommend that APM clients use round-robin access from the host. Doing so ensures that all ports and Pillar Axiom CPUs are fully utilized.

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

For release information, 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.

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 5 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 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 Pillar Support (http://www.pillardata.com/support).
2. Click **CUSTOMER LOGIN.**
3. Enter your username and password.
4. In the left navigation pane, click **Software Downloads**.
5. In the left navigation pane, click the platform or software category in which you are interested.

6. In the content pane, navigate to and click on the software package you want. Download details for that software appears in the bottom of the content pane.
7. Click the green arrow icon or the **Click here to download SW** link.

Note: To obtain a license to enable a feature that is in addition to those that you initially purchased, contact a Pillar sales representative. (See Table 5 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 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 http://www.pillardata.com/techdocs.

5 Pillar Axiom System Limits

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

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

- You configure a system near one or more of its limits.
- You increase the number of customer-defined system objects — 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 6 Operating limits for ALL Pillar Axiom systems

Item	Description and range
Volume groups	Minimum = 1 Maximum = • 5000 total, out to five levels • 100 subgroups for a given volume group
Repository VLUNs [f]	Maximum = 1024

Table 7 Operating limits for Pillar Axiom NAS systems

Item	Description and range
File Servers [g, h]	Maximum = • 4 for each Pillar Axiom 300 Slammer • 8 for each Pillar Axiom 500, 600, or 600e Slammer
VIFs for each File Server	Minimum = 1 Maximum = 16
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
Static and default network routes for each File Server	Minimum = 0 Maximum = 8 default, 16 static
NIS alternative file size	Up to 50 MB

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

[g] The virtual interfaces (VIFs) of a File Server can be configured on multiple Slammer control units (CUs). 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. Pillar recommends, however, that multiple File Servers be defined across a collection of Slammer CUs rather than spreading a single File Server across those CUs.

[h] 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.

Item	Description and range
Filesystems	Minimum = 1 Maximum[i] = 1024
Filesystem size	Minimum: 1 to 2 GB. The exact value depends on these factors • Brick type (Fibre Channel, SATA, or SSD) • RAID geometry (RAID 5 or Pooled RAID 10) • Strip size (1 MB or normal) Maximum[j] = system capacity
Filesystem name length[k]	Maximum = 63 bytes
Directory quotas	Unlimited
Snap FSs	Maximum = • 250 for each filesystem • 16,000 for each Pillar Axiom system
Clone FSs (partial block snapshots)	Maximum[l] = 4095 for each Pillar Axiom system
Replication targets	Maximum = 5 for each replicated filesystem
Concurrent NDMP sessions	Maximum = 10 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[m]	Maximum for each Slammer (combined memory for both control units): • 400 for 6 GB combined memory (Pillar Axiom 300 systems only) • 1200 for 12 GB combined memory (Pillar Axiom 500/600 systems only) • 6000 for 24 GB combined memory (Pillar Axiom 500/600 systems only) • 12000 for 48 GB combined memory (Pillar Axiom 600 systems only)

[i] The maximum number of filesystems applies both to the system and to a NAS Slammer.

[j] Maximum capacity is *unlimited* for a logical volume if the Thin Provisioning license is installed.

[k] Maximum name length applies to Snap FSs and Clone FSs as well. When using 3-byte UTF-8 characters, the limit is 21 characters.

[l] The 4095 maximum is the maximum number of VLUNs supported in a Pillar Axiom system. VLUN structures underlie all filesystems, LUNs, Clone FSs, and Clone LUNs. This 4095 maximum covers all four volume types, in total.

Table 8 Operating limits for Pillar Axiom SAN systems

Item	Description and range
SAN LUNs	Maximum [n] = • 4095 visible for each system • 256 visible for each host • 4095 visible for each SAN Slammer
SAN LUN size	Minimum: 1 to 2 GB. The exact value depends on these factors • Brick type (Fibre Channel, SATA, or SSD) • RAID geometry (RAID 5 or Pooled RAID 10) • Strip size (1 MB or normal) Maximum [o] = system capacity
Volume Copies (full block snapshots)	Maximum = • 1024 for each Pillar Axiom system • 12 active for each LUN
Replication pairs	Maximum = 16 for each Pillar Axiom system
Clone LUNs (partial block snapshots)	Maximum [p] = number of unallocated SAN LUNs, up to 4095
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

[m] The maximum number of CIFS connections supported in failover mode is the same as the number supported on a single CU, which is ½ of the maximum for the entire Slammer. However, for 24 GB and 48 GB Slammers, Pillar highly recommends a maximum of 2400 CIFS connections (1200 for each CU).

[n] See footnote *l* for more information regarding this maximum.

[o] See footnote *j* for more information regarding this capacity maximum.

[p] See footnote *l* and footnote *g* for more information regarding this maximum.

5.2 Slammer and Brick Configuration Limits

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

Table 9 Brick configuration limits

Number of Slammers	Minimum number of Bricks		Maximum number of Brick strings	Maximum number of Bricks
	Supported	Recommended		
1 [s]	2	3	4	32
2	4	All NAS or all SAN Slammers = 5 NAS / SAN combo Slammers = 6	8	64
3	6	8	8	64
4	8	8	16	64

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

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

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

For Fibre Channel (FC) Bricks:

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

  Note: Version 2 FC Bricks do not have an FC Expansion Bricks option.
- Pillar Axiom systems have a limit of 32 FC Bricks, regardless how many Slammers are in the system.

[q] The minimum configuration for Pillar Axiom 300 systems is one Slammer and one Brick (SATA or FC). The maximum number of Bricks in Pillar Axiom 300 systems is eight.

[r] The minimum configuration for Pillar Axiom 600e systems is one Slammer and one SATA or FC RAID Brick (FC Expansion and SSD Bricks are not supported). The maximum number of Bricks in a Pillar Axiom 600e system is two.

[s] For single-Slammer Pillar Axiom 500 configurations and single-Slammer Pillar Axiom 600 configurations with the Extended Brick license installed, the minimum number of Bricks is two of the same type. However, for mixed configurations, the minimum number of Bricks is three, as outlined below:

- For a mix of FC and SATA (or SSD) Bricks: 2 SATA (or SSD) + 1 FC or 2 FC + 1 SATA (or SSD).
- For a mix of SSD and SATA Bricks: 2 SATA + 1 SSD or 2 SSD + 1 SATA.

- FC Bricks may co-exist with SATA and SSD Bricks in the same string, subject to current configuration recommendations.
- A given Brick string can contain up to four FC Bricks (RAID or Expansion). Maximum number of FC Expansion Bricks in any string, however, is two.

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

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.

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

6.2 Network Requirements

6.2.1 Pilot Network Requirements

The Pilot management controller requires:

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

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

The 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 1 Gb/s or 10 Gb/s connections. Both fiber and copper are supported.

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

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

6.3 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 robotic library software 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 an emergency, drop all power and then contact the Support Center. Contact the Support Center before touching any power cables or switches. There are some situations where *not* power cycling the entire system is the correct action.

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

See also Sections 9.57, 9.62, 9.84, and 9.98.

7 Known Issues

The Pillar Axiom server issues listed in Table 10 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 5 Contact information).

Table 10 Known Pillar Axiom server issues

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
All	Multiple concurrent restore operations on the same volume from a clone of that volume do not work. The first restore may succeed, but the second hangs and be outstanding forever. The situation gets more complicated if a Slammer control unit fails while the restore operations are in progress.	Avoid running concurrent restore operations for a volume. This issue will be fixed in a future release.
	Updating software from the AxiomONE Storage Systems Manager (GUI) sometimes fails with a “Software update package is invalid” error, and an Administrator Action may appear indicating failure of the PerformUpdatePackageStagingTask.	Use the following pdscli command instead: `pdscli sub -u <username> -p <password> PerformUpdatePackageStaging FileContents="&@<path-to-package>"` This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	SSH access to the Pillar Axiom system is useful to aid in problem diagnosis and resolution. However, if the Support user has not previously been granted the "enable SSH" capability, it is not possible to grant this capability after the malfunction has happened, because modification of Account attributes requires that the system be in a Normal state.	No workaround once the malfunction occurs. To enable SSH access when there is a problem, turn on Enable SSH capability in the Support account while the Pillar Axiom system is functioning normally. Note: Turning on the Enable SSH capability does not enable SSH. It only allows the Support account to enable SSH. This issue will be fixed in a future release.
	Under a specific heavy workload pattern, a Brick RAID controller can sometimes run out of resources, resulting in I/O to that controller being stuck for a certain amount of time. This can result in pinned data and the associated LUNs alternating between Online and Conservative.	Contact Pillar Worldwide Customer Support Center for help in resetting the RAID controller. This issue will be fixed in a future release.
	Resetting the Support Administrator credentials by means of the USB key reset procedure is not supported.	This issue will be fixed in a future release.
	The Pilot may fail to restart following a power outage if the Pilot was writing to the drive at the time the power failure occurred. This is caused by a single, uncorrectable sector on the Pilot drive and the following message appears: "fsck.ext3: Input/output error while recovering ext3 journal of /var."	Restart the Pilot. This issue will be fixed in a future release.
	During very short (<200 msec) AC power disruptions, a Slammer control unit (CU) can sometimes go offline and stay offline.	Power cycle the Slammer CU. Note: Data in battery-backed memory is kept intact during the power cycle. This issue will be fixed in a future release.
	When a Brick experiences a read error on two drives in the same stripe or a read error during a RAID array rebuild, the data for the associated stripe is lost and a Bad Block Log entry is created to track the bad stripe. If any reads are sent to this stripe, they are rejected with an error. If the stripe is written, the Bad Block Log entry will be cleared.	Contact the Pillar World Wide Customer Support Center for assistance. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	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 9.97 for more information.)* An upgrade to the latest release of OpenSSH will be available in a future release.
	If a system software update is initiated while Pilot tasks are in progress (such as a SecureWORMfs scan, automatic Call Home, 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.
	When a Brick RAID controller is removed during Guided Maintenance, the remaining controller may send a “RAID Controller Fault” event to the host in addition to the “RAID Controller Removed” event.	This issue will be fixed in a future release.
	On rare occasions when upgrading system software, the firmware upgrade in a Brick RAID controller can fail to complete, which causes the system upgrade to fail.	Replace the RAID controller that has failed the firmware upgrade. Then restart the system software upgrade. This issue will be fixed in a future release.
	In rare situations, when a SATA RAID controller is failing and a drive in the same Brick is taking a very long time to execute commands, a RAID controller reset may take the whole Brick offline.	Address the RAID controller failure and the drive issue separately. This issue will be fixed in a future release.
	Call Home transfers by means of an HTTPS proxy server may be reported as successful on the Pillar Axiom system when in fact the Call Home log bundle has not been received by Pillar.	Verify whether the system serial number of the Pillar Axiom system is properly registered with the Pillar Call Home server so the proxy server can login and send Call Home log bundles. This issue will be fixed in a future release.

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.
	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.
	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.
	Pillar Axiom systems do not support the use of Windows 2003 SP1 or higher as the NTP server for the system.	For an explanation of how to provide a stable NTP for both Windows and Unix environments, refer to Microsoft TechNet article "Appendix H: Configuring Time Services for a Heterogeneous UNIX and Windows Environment" (http://technet.microsoft.com/en-us/library/bb463171.aspx). The Meinberg NTP server is a software product that can be installed on a Windows platform to provide NTP services for a Pillar Axiom system. This issue will be fixed in a future release.
	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 drive is performing marginally, the system does not notify the storage administrator until the error-rate becomes high enough that the drive is taken offline. At that point, the spare drive is used automatically.	This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	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	Performing an NDMP file-based restore to a full filesystem can fail even when the restore is overwriting the original files. This happens when large files are deleted in the overwrite process. The large files are deleted in the background while the NDMP restore is writing new files in parallel. If the delete is slower than the restore, the filesystem may temporarily run out of space.	Here are two alternatives: • Restore to a filesystem with enough space to hold the restored files. • Grow the filesytem being restored from the backup. This issue will be fixed in a future release.
	From a Windows 7 client, one might not be able to delete a folder that is on a share on a Pillar Axiom File Server if the Read Only Attribute is set.	This issue will be fixed in a future release.
	From a Windows 7 or Windows 2008 client, attempting to change the security permission of a CIFS share on a Pillar Axiom system can result in an "access denied" error. This problem does not occur on Windows XP client.	This issue will be fixed in a future release.
	In rare circumstances, an internal filesystem named PDS might not automatically be recreated when it has a data loss, which may lead to complete or partial loss of NAS services until the filesystem is made whole again.	Contact Pillar World Wide Customer Support for help in recreating the filesystem. This issue will be fixed in a future release.
	Under very rare circumstances, the Pillar Axiom system can warmstart when there are two simultaneous internal requests to convert DNS hostnames or IP addresses to internal data structures.	This issue will be fixed in a future release.
	When configuring for link aggregation, if the Slammer ports are configured first and then the switch, the client will appear to lose its connection to the ports.	Disable and then re-enable link aggregation on the Slammer ports. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Sometimes the information related to the CIFS shares displayed by the GUI can be incorrect when the owning Slammer control unit warmstarts before the persistence database is updated after certain operations related to shares.	This issue will be fixed in a future release.
	Intermittent warmstarts can sometimes occur when deleting shares while CIFS clients are active.	If a warmstart occurs, perform these actions: 1. Stop CIFS access (or disconnect the data network). 2. Delete the shares. 3. Resume data access (or reconnect the data network). This issue will be fixed in a future release.
	Changes made in the Shares tab on the **Storage** > **NAS** > **Filesystems** > **Modify Filesystem** page in the AxiomONE Storage Systems Manager (GUI) sometimes does not take effect.	Use the **Storage**> **NAS** > **Shares** page instead. This issue will be fixed in a future release.
	NDMP file-based backup software sometimes encounters an error during the encoding process when running an NDMP backup, which causes the Slammer control unit (CU) to warmstart.	Wait for the Slammer to complete the warmstart and restart the backup when possible. This issue will be fixed in a future release.
	Changing the virtual interface (VIF) settings for a File Server in the AxiomONE Storage Systems Manager (GUI)—for example, to use the same IP address but a different gateway or netmask—sometimes causes the modify operation to fail with an error code of 6504.	To change the VIF settings: 1. Delete the existing VIF. 2. Create a replacement VIF with the desired properties. This issue will be fixed in a future release
	Leaving the **Default Gateway** field blank in the AxiomONE Storage Services Manager (GUI) Create File Server page causes a failure. Leaving **Default Gateway** blank in the Modify File Server page leaves the File Server in an inconsistent state.	Always provide a default gateway value when creating or modifying a File Server. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	In the GUI, the **Identify Port** button function on the Modify File Server page, which is used for blinking LEDs to identify associated ports, is not supported on 10 GbE NIMs.	This issue will be fixed in a future release.
	On occasion, when I/O is occurring, it's possible that something can go wrong with a spacemap block causing the system to quarantine the block. When this occurs, the system also creates an Administrator Action that states "FileSystem SpaceMap Block Quarantined" followed by the filesystem name.	Spacemap quarantines do not cause any data or metadata loss or corruption. The system should remain in a "green" state while production I/O continues. This issue will be fixed in a future release.
	NDMP and NAS Replication operations fail if the NDMP Backup Settings in the AxiomONE Storage Services Manager (GUI) include a File Server with spaces in its name. These operations will continue to fail even after the spaces have been removed from the File Server name.	1. On the Modify File Server page, change the NDMP File Server name to a name with a single space in it. The part of the name after the space must include at least 8 characters, as in "fserv fixissue." 2. On the NDMP Backup Settings page, do the following: a. Set the File Server to "None." b. Set the File Server name to the name with a single space in it, as in "fserv fixissue." 3. On the Modify File Server page, reset the File Server name to the desired name and ensure that it has no spaces. 4. On the NDMP Backup Settings page, do the following: a. Set the File Server name to "None". b. Reset the File Server name to the desired File Server name. This issue will be fixed in a future release.
	Under rare conditions, even though the user interface lists a particular share as existing on the filesystem, that share might not actually exist. Also, attempts to recreate the "missing" share can fail.	Restart the Pillar Axiom system. If a system restart is not desirable, contact the Pillar World Wide Customer Support Center for assistance. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	When creating an NFS export using the AxiomONE CLI (axiomcli) and you want to specify the user ID to use for anonymous access (-anonuid), you cannot use "0".	Use a value other than "0" for the anonymous user ID. If the user ID for anonymous access must be set to "0", use the pdscli instead. This issue will be fixed in a future release.
	A NAS Replication `set_live` command fails if a full synchronization of the filesystem has not completed. When this occurs, sometimes the error message is misleading.	If the GUI lists the filesystem on which the `set_live` command failed as "Prepared for Restore," the failure occurred because the filesystem has not been fully synced. Fully sync the filesystem and re-issue the `set_live` command. This issue will be fixed in a future release.
	The user filename parameters of the axiomcli or pdscli requests for uploading NIS files can only be filenames. If they are pathnames to the files, the requests will fail.	Place the files in the same directory where the upload script is located. Then, use only the file names to identify the files to be uploaded. This issue will be fixed in a future release.
	Issuing a `discover` command on a NAS replication pair may cause a currently running synchronization operation to fail.	Avoid running the `discover` command on a replication pair in this situation. Instead, monitor the output from the `sync` command. Resubmit the `sync` command, which will cause synchronization to restart from the point of failure. This issue will be fixed in a future release.
	Under certain conditions, when the amount of resources required to maintain a large number of open files (e.g., 1200 files) or a number of files that have path names of 100 or more characters, plus the resources for the other activity on the connection cause an internal threshold to be exceeded, Kerberos authentication requests or file open requests may be denied. In some case, the Slammer control unit may warmstart.	In these scenarios, each user should use a separate connection by running one user or application for each Windows client (or for each virtual client with a different IP address). Contact Pillar World Wide Customer Support Center if you need additional help. This issue will be fixed in a future release.
	Cannot delete a directory quota on a non-empty directory while the filesystem is online.	To delete a directory quota on a non-empty directory, first take the filesystem offline. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Sometimes the event for "Quota hard limit" may not be generated. Though the event is not generated, the hard limit is honored.	This issue will be fixed in a future release.
	When CREATOR_OWNER role is expanded, the newly created ACL entry (ACE) does not appear as inherited.	This issue will be fixed in a future release.
	If there are many disabled server entries listed in the DNS, the Pillar Axiom system may fail to join the domain due to no response from them and a time out (250 seconds).	Remove the disabled Domain Controllers from the domain. This issue will be fixed in a future release.
	When doing directory listings using OpenDir API, the directory entries may not always fit into the system memory because of its size. In the case where the memory size did not fit with the directory listing, the operation may fail.	This issue will be fixed in a future release.
	When the DNS server is slow and multiple local Domain Controllers are down, the Pillar Axiom system may not be able to join the domain using a remote Domain Controller if the time remaining allowed to join the domain expires.	Ensure the DNS server points to functional and working local Domain Controllers. 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.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	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.
	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.
	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.
	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 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.
	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.
	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.
	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.
	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.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
SAN	If an iSNS Server has been registered on a Pillar Axiom system and the iSNS Server cannot be contacted by either static IP address or DHCP, the system may not be able to fully shut down. If a shutdown operation is requested, the task completion operation may stop at less than 100%. This can prevent the system from performing major release software upgrades.	Ensure that the Pillar Axiom system can contact the iSNS Server before initiating a system shutdown or a major release software upgrade. Alternatively, temporarily disable iSNS registration on the Pillar Axiom system before initiating a system shutdown or a major release software upgrade. This issue will be fixed in a future release.
	The Brocade HBA may not function properly when connected directly to the Pillar Axiom system. The HBA appears to be connected; however, no LUNs are visible on the host.	Connect the Brocade HBA using a switch or use loop mode. This issue will be fixed in a future release.
	When using a Fibre Channel HBA in a Windows environment, if a SAN error is encountered, recovery may take up to four minutes and can result in performance issues.	Inspect your SAN configuration for bad connections and resolve any issues that may exist. This issue will be fixed in a future release.
	The task to remove a Brick may fail if a system has LUNs that participate in replication pairs, even though all of the volumes that claim to be using that Brick have been deleted.	Stop all replication and delete the replication pairs. Attempt to remove the Brick again. This issue will be fixed in a future release.
	On the Identity tab in the GUI, when activating an inactive Clone LUN that was created by SAN Replication, if any other tab is visited before un-checking the Inactive checkbox, the protocol choices for LUN Access are cleared of previous values. Also, an error occurs and the user needs to return to the Identity tab and specify the desired protocols.	Un-check the Inactive checkbox before leaving the Identity tab. This issue will be fixed in a future release.
	A defective Fibre Channel network interface module (NIM) may result in a kernel panic and a warmstart of the Slammer control unit.	Replace the defective NIM. This issue will be fixed in a future release.
	For a given host system, only one axmrepsan session can be active for a given unique userid. Attempts to run more than one axmrepsan session for a given userid will fail with unpredictable results.	Have a unique userid for each axmrepsan session. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	Continual SNMP querying of "cSANHostConfig" causes a slow but steady consumption of Pilot CPU and memory resources, causing degradation in Pilot performance and ultimately Pilot failover.	Avoid SNMP queries for this item, or do so as infrequently as possible. 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.
iSCSI	In rare circumstances, iSCSI systems under heavy load or with high numbers of network errors may warmstart.	Try to rebalance the load. If you're getting a large number of network errors, replace the iSCSI network interface module in the Slammer. This issue will be resolved in a future release.
	Modification of iSCSI port settings using the `ModifyiSCSIPortDetails` CLI request fails if attempted within two minutes of a system restart.	Wait at least 2 minutes after a system restart completes before executing a `ModifyiSCSIPortDetails` request. This issue will be resolved in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	On rare occasions, the iSCSI network interface card can hang, which causes the system to warm start.	Replace the iSCSI card, or contact Pillar World Wide Customer Support for assistance. This issue will be fixed in a future release.
	During Windows startup, the iSCSI Software Initiator may attempt to register with the Microsoft iSNS Server v3.0 if it has been configured to do so. If the iSNS Server is installed on the same host and has not started yet, the iSCSI Initiator may fail to register with the server.	Edit the Windows registry to add a dependency that causes the iSCSI Initiator to wait for the iSNS Server to start. This issue will be fixed in a future release.

8 Resolved Issues

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

Table 11 Resolved Pillar Axiom server issues

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
4.3	All	During a warmstart, attempts to reinitialize the internal Ethernet controller can sometimes fail, which results in a Slammer control unit (CU) failover.
		Unable to parse and see Brick S.M.A.R.T. data using the Pillar AxiomONE Statistics Tools.
		In rare cases, the core dump may be unavailable after the Slammer warmstarts.
		Relevant Brick logs might not be automatically collected for Call Home when a fault event for a RAID controller occurs.
	NAS	While running concurrent full NAS replications at relatively slow speeds, the target Pillar Axiom system might warmstart. This could happen, for example, when setting up two Pillar Axiom systems and performing initial synchronization operations to the target over a WAN link.
		When a Slammer control unit (CU) fails over, the buddy CU takes ownership of each filesystem owned by the failed CU. For each filesystem, the I/O is quiesced and the filesystem remounted on the buddy. Due to backend issues, the remount could not be completed timely and was retried, resulting in the internal health checker timing out the quiesced I/O. This timeout results in the CU being reset.
		Pinging a File Server virtual interface (VIF) might fail if one or more ports in the NAS Slammer control unit (CU) do not have physical connectivity. The ping might fail even if the VIF is not associated to the failed port. For example, in a 4-port system, if a VIF is configured on port3, but port2 (its failover partner) is physically disconnected, pings to the VIF might fail. New network traffic to the VIF might not be successful; existing connections, however, will continue to work.
		When a NIS server sends a record of length greater than 1024 bytes, the Pillar Axiom system stalls and a software fault occurs. During this time authentications will fail.
		When an NFS mount operation fails, there is no event containing the original path supplied by a client.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		The Pillar Axiom system sometimes cannot detect all tape devices when the external Fiber Channel (FC) switch that connects the FC tape HBA in the Slammer to the target tape devices is in the same zone as other initiators.
		When a VLUN is offline and takes a long period of time to recover, the filesystem may remain offline despite the VLUN being online. In those cases, a supporttool command to manual restart with battery-backed memory preserved is needed to bring the filesystem back online.
		It's not possible to clear the filesystem check (FSCK) logs from a NAS system. This eventually causes problems with log collection as the logs grow in size.
		Sometimes a large file having a large number of asynchronous writes, in which some of those writes have already been committed to the journal but not properly accounted for as committed, can cause a software fault.
		User account mapping is supposed to map a CIFS account to a combined identity consisting of the user's CIFS identity and NFS account. The combined identity, however, was only associating one of the user's NFS groups.
		The system generated 90-day validation scan on WORM filesystems impacts production I/O.The fix removes the 90-day validation scan. However, to manage its impact in a more flexible way, an administrator can manually perform a validation scan by choosing **Perform Protected File Integrity Scan** from the Actions drop-down list in the GUI.
		During the recovery period of the power-on-data-recovery (PODR) process, if the VLUN of a filesystem is offline, the filesystem skips the replay of journal of the filesystem and marks the filesystem offline. The journal of this filesystem, if any, can be inadvertently overwritten by other filesystems, leading to filesystem corruption. Also when the vLUN is back online, the filesystem is not brought back online and requires a PODR to bring the filesystem online.
		During the recovery period of a filesystem following a power cycle, the filesystem may be marked Offline with no actual journal recovery performed if the underlying physical storage is not yet available. In this Offline state, the journal associated with this filesystem might be overwritten by other filesystems, causing subsequent data corruption.
		The local users and local groups are not included sometimes in the enumerated list for selection when modifying file and directory permissions.
		When the filesystem is in a degraded condition due to control unit (CU) failure, the filesystem relies on disk based journaling. In rare circumstances, when the surviving CU crashes during recovery of the disk based journal, part of the journal update could be missed, which leads to filesystem corruption.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		When the system detects a zero reference count in a used superblock buffer, the system places a “PANIC_ASSERT: DJC: Bad refcount” message in the system log and warmstarts the Slammer.
		To better interoperate with switches that support SMLT (Split Multi-LinkTrunking), Axiom should not set its System Priority for switches to levels lower than the switch’s own priority. The Nortel switch sets its default priority to 32768. The fix changes the system priority on the Pillar Axiom system to 60000.
		The local Administrator and local Backup Operator groups are not included in the enumerated list for selection when modifying file and directory permissions.
		When creating an NFS export using the AxiomONE CLI (axiomcli), specifying "0" for anonuid fails with the following error: Error: Please provide a user id for anonymous access.
		When the I/O from the backend is slow, some filesystems can run out of journal space in the battery-backed memory, causing these filesystems to switch to a disk-based journal. However, when the issue is resolved, the filesystem fails to switch back to the normal mode of using the memory-based journal.
		When more than 1023 customer filesystems are created, the Slammer is not able to successfully handle failover, which results in both control units being disabled.
		If a filesystem is renamed while it is offline because of corruption, sometimes the system might erroneously put the filesystem online. In this case, the filesystem will go offline again when the system detects the corruption.
		When a Pillar Axiom system undergoes warmstart recovery or upgrade, the underlying I/O subsystem may become congested. If the congestion lasts more than five minutes, the filesystem may go offline and have to be manually brought online.
		When a filesystem is corrupt, an attempt to create clone space for this filesystem will fail, generating a Failed ModifyFileSystemTask Administrator Action.
		After an initial query, some filesystem statistics on subsequent queries return 0 in the output of the GetFileSystemDetails CLI command and in the display on the Health > Performance > Filesystem GUI page. This occurs, for example, for the "ReadMBPerSecond" and "Current MB per second" statistics in the CLI and GUI, respectively.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		A filesystem that cannot flush its journal to disk generates a pinned data condition. When a filesystem that is being deleted has a pinned data condition, the system automatically discards the pinned data. This results in a BBM (battery-backed memory) lost event.
		Sometimes, when a system is under heavy I/O loads, the system may slow and Slammers may begin to warmstart because of I/O timeouts. Also, under these conditions, a filesystem may show a Conservative status.
		The system may incorrectly 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.
		When using a device with less than 1 Gb/s link speed to connect to a NAS Slammer, the I/O Port Details screen will indicate 0 Gb/s for the Negotiated Link Speed
	SAN	When isolating a replication pair using the command `axmrepsan isolate_replication_pair` ***`pair_name`***, sometimes the replication pair can go into the AWRY state, which is verified in the output generated by the command `axmrepsan get_replication_pair_status`.
		If there were a large number of SAN cache de-stage requests pending, an incoming write from a host could be delayed long enough that the host would abort.
		Due to a software race condition, SAN LUNs could fail to come back online after a system power failure.
4.2	All	Following the addition of Bricks and code upgrades, sometimes the system inadvertently turns off Brick event monitoring. As a result, the system is not aware of problems and of error recovery actions being taken by the Bricks.
		When a Slammer control unit failback fails and a second failback succeeds, the battery-backed write cache may not be properly established for logical volumes, causing a significant performance loss.
		Due to an internal locking issue the Pillar Axiom system may warmstart in overloaded conditions.
		When an external security audit is run against the Pilot management controller, sometimes the audit indicates a security violation (the server supports the TRACE and/or TRACK methods).
		A failing drive on a Fibre Channel Brick may cause the Brick to become unresponsive to I/O requests for a few minutes if a write command fails to the drive. This lack of response may cause LUNs and filesystems to go offline temporarily.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		When a Pillar Axiom system experiences several Slammer control unit failovers in a short period of time, sometimes the system stays in write-through mode after failback, causing performance to degrade severely.
		A Slammer control unit may warmstart unexpectedly after performing management operations (such as creating and deleting Clone LUNs) for several hours.
		After a successful shutdown, sometimes a restart can result in the system erroneously reporting lost blocks.
		Pillar Axiom 300 systems may not be able to complete a warmstart successfully, leading to a Slammer control unit failover or failback when the warmstart times out.
		Discovery of a triply redundant volume does not produce a configuration-on-disk (COD) error as expected.
		When a Pilot control unit (CU) is replaced in a Pillar Axiom system, sometimes the active Pilot CU fails to update some of the components on the replacement Pilot CU correctly.
		Events are generated for non-critical internal I2C bus errors.
		Not able to download the Pillar Statistics Tools for Windows and Linux platforms.
		During a Slammer failover operation, an additional Slammer warmstart may result.
		System failure may result from a user application attempting to write to a volume that is in the process of being thinly provisioned.
		If a Slammer control unit (CU) warmstarts while the system configuration is being changed, the CU may subsequently warmstart again.
		If one creates a thinly-provisioned volume, and then makes it fully provisioned by changing the current capacity to equal the maximum capacity, sometimes an error can occur (because of internal rounding errors).
		A memory leak in the message buffer may cause a Slammer control unit to warmstart unexpectedly during routine management operations that last several hours.
		When deleting an extremely large LUN or filesystem, the Brick takes a large amount of time to complete the transition of the related storage from allocated to unallocated. During this time, the Brick is temporarily inaccessible until the operation completes.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		During failover and failback, a Slammer control unit may warmstart unexpectedly.
		A Slammer control unit (CU) may fail when intense I/O on a sparse volume causes a high rate of infill allocations in a short time. This high rate of allocations can cause a count mismatch, resulting in CU failure.
		During a software update, you may receive an Administrator Action indicating that the Update Software Task Failed. This failure may be the result of some internal events that may have not been forwarded (or processed) during a Slammer control unit warmstart.
		PIM replacement, which includes powering off and powering on of the parent Slammer control unit (CU), may cause NonOptimizedAccess events for APM.
		The `GetStorageConfigDetails` PDSCLI command returns raw capacities based on a RAID 5 geometry. This does not account for capacity requirements when using Pooled RAID 10.
	NAS	For NAS Replication, running 10 parallel sync operations with the "diffsync=yes" flag set can sometimes result in failure when attempting to start the 10th sync operation. This condition persists until a sync operation that is currently running completes.
		For some authentication to work, a symbolic link is needed to access local files. The actual making of the link on a warmstart was lost and access to the Pillar Axiom filesystem was not available.
		When a Windows 7 client uses a roaming profile configured in Windows Server 2008 to create a home directory on a Pillar Axiom system, the operation appears to complete successfully. However, when the user logs off, the user settings are not retained and the user receives the following error: "The user profile was not completely synchronized. Please read event logs for details."
		When a directory is very restrictive, the security attributes may make it impossible for a CIFS administrator to do anything to the directory, even taking ownership of the directory. If the directory is the root directory of a filesystem, the entire filesystem can become unusable.
		When the record count of a CIFS share table is incorrect, the CIFS server panics, resulting in a Slammer warmstart. The fix causes the CIFS server to delete the corrupted share table, resulting in the loss of the CIFS shares. Administrators will need to recreate those shares.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		When a filesystem with locks is re-homed and the NFS client that has the locks either reboots or crashes, the NFS client (when it comes back up) will issue a NLM FREE_ALL request, which causes the Slammer CU to warmstart.
		Changing the configuration of Windows domain controller without a warmstart may cause the CIFS server to consume extra memory to store the new setting. Over time, the CIFS server runs out memory and fails to create connection to Windows domain controller for authentication.
		On rare occasions, a NAS replication target filesystem can become corrupted when running multiple replications on the same replication pair at the same time.
		When you move a subdirectory to a different parent directory and then delete the former parent, if you then create a new file (fileX) somewhere and create a hard link (hdlinkK) to that file from the above subdirectory, attempting to move another file into that subdirectory results in an error.
		Sometimes in the AxiomONE Storage Services Manager, you can perform several separate, but related, actions in a single, overall operation. For example, you can create a File Server, a filesystem, and an NFS Export as part of a single create operation. In these situations, if one of these actions cannot be done for some reason, all of the actions cannot complete and the recovery and rollback of the actions may not be done correctly.
		When a NAS control unit (CU) fails over, the system reassigns the resources to the buddy CU. However, under rare circumstances, the CU may fail a second time.
		When using InMage for replication with a Pillar Axiom system as the target, the file attributes have the HIDDEN attributes set.
	SAN	During replication, a background data migration deadlock can occur that causes a Pillar Axiom system to warmstart.
		On rare occasions, the Pillar Axiom system has problems tracking SAN host logouts and logins, which can result in a warmstart of a Slammer control unit.
		Migrating a currently replicating LUN to another Storage Class may result in a system warmstart.
		Unable to save a change only to the access bias of a LUN. Additional changes to the LUN were required for the save to be successful.
		The SAN Remote Replication feature requires both a Remote Replication license and a Clone LUN license to function. The fix requires only the Remote Replication license.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		Unexpected errors may occur when attempting to move LUNs between volume groups. The system may think that there is insufficient capacity in the volume group for which to move the LUN. Also, the reporting of these errors may state that it is an "Xpath error" which can be ignored.
		During the checkpoint process, a SAN Replication pair may go AWRY.
		When SAN Replication pairs are running in two directions at the same time (source to destination and destination back to the same source), the pairs may stop replicating, eventually causing the pairs to go AWRY or ISOLATED.
		When SAN Replication pairs are running in two directions at the same time (source to destination and destination back to the same source), the pairs may stop replicating, eventually causing the pairs to go AWRY or ISOLATED.
		During replication, the Pillar Axiom system might experience multiple software faults, leading to a full system restart.
		When you create a SAN Replication checkpoint for which the source is ready but the target is not, numerous Administrator Actions stating that the target checkpoint failed are issued.
		When the primary and secondary Pillar Axiom systems hosting a replication pair are started at the same time, sometimes one or both systems may not complete the restart, which can result in Pilot failover and a significant delay in getting both systems operational.
		During SAN replication operations, if failover-failback occurs, additional warmstarts might occur, disabling the Slammer control unit.
		In a clustered environment that uses SCSI persistent reservations, a SAN host that has been rebooted and lost its persistent reservation or registration for a cluster LUN (which had LUN mapping enabled) can still access the LUN.
		After recovery from a Slammer control unit (CU) failure, the Pillar Axiom system may become incapable of automatically moving LUNs between the CUs on that Slammer. When the system attempts to move the LUNs automatically in response to non-optimized access from a host, the attempts fail, and non-optimized access persists.
		When AxiomONE Replication for SAN replication pairs are replicating, numerous Event Out Of Sequence events can appear in the event log.
		When the owning Slammer of a Clone LUN is in a non-normal status, the health of the Clone LUN may incorrectly be displayed as Online when in fact it is Conservative.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
4.1.1	All	When deleting a replication pair, the Pillar Axiom system may experience failures leading to an offline state.
		During an upgrade to release 4.x, it is possible for uninitialized legacy Array Manager metadata to be erroneously processed, leading to the potential for a false declaration of lost blocks.
		A software fault may occur during replication due to a flaw in data flow control.
		When modifying both the storage class and priority of volume (LUN or filesystem) the new priority setting was used as a temporary value instead of a fixed value. This only happened if the selected value in the storage class dropdown menu was changed before the selected value in the priority dropdown menu was changed.
		There was a rare Fibre Channel Loop occurrence that could lead to the Brick firmware falsely raising a particular error condition. This would result in Brick drives being taken offline.
		Modifications to either the primary DNS server or the secondary DNS server or both as part of the Call Home configuration did not save correctly.
		In the event of a RAID controller debug restart, it was possible that the debug output would not be properly extracted.
		When a Pilot control unit (CU) is replaced in a Pillar Axiom system, it is typically updated with the proper software by the primary Pilot CU. There were cases, however, for Pillar Axiom 600 systems, when the active Pilot CU would fail to update some items on the replaced CU correctly.
		When a Clone LUN or Clone FS is deleted from the middle of a hierarchy of such clones, the system could get confused about the relationships of the remaining clones.
		If the Fibre Channel interconnect between the Slammer and Brick was experiencing a large number of LIPs and aborts it was possible for the Brick to run out of transaction resources resulting in it being unresponsive to I/O.
		When deleting large filesystems or LUNs, the system could experience a warmstart.
	NAS	When modifying the default quota limits, the individual user quota limits did not change.
		VLUN metadata could be improperly handled during code level upgrades, resulting in the system erroneously reporting Lost Blocks.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
	SAN	The ATIME (last accessed) attribute was not being set on NAS files by default. This fix changes the default behavior for new filesystems to update ATIME.
		Users could not view used repository capacity without a SnapLUN license, but users of RMRLite needed to be able to do this without a SnapLUN license present.
		Incorrect requirement for replication licenses on both the primary and secondary sites in order to do replication. With this fix, replication license is only required on the primary side.
		The primary and secondary replication points represented different capacity values to the host.
		When issuing a replication checkpoint where the source was ready, but the target was not, users would receive an Administrator Action that the target checkpoint had failed.
		During replication rebuild, if the operation needed to be restarted, it was possible that the rebuild operation would not start.
		During replication, there was a small chance that the Pillar Axiom system would experience multiple software faults, leading to a full system restart.
		In some situations, a Clone LUN that had been previously deleted could reappear after a system power cycle.
		Guided Maintenance for PIM replacements, which include powering off and powering on of the parent Slammer control unit (CU), may cause NonOptimizedAccess events for APM
		If a LUN were modified by changing only the access bias, the change would not save correctly.

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 system, please refer to the *Pillar Axiom Administrator's Guide* for more information.

9.1 Host Queue Depth on SAN Hosts

The recommended maximum queue depth for all SAN hosts attached to a Pillar Axiom system is 64. This value is the maximum number of outstanding I/O requests to the Pillar Axiom system. Exceeding this value may cause I/O errors if the input/output queue of the Pillar Axiom system is exceeded.

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

9.2 Limitation of the Linux sginfo Utility

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

9.3 LUN Ranges in Windows

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

9.4 Issues with LUN Capacity Calculations on Solaris

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

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

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

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

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

9.5 VMware ESX Server 3.0.1 Connections to Pillar Axiom iSCSI Systems

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

9.6 Avoid 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

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

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

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

9.9 LUN Assignment and Accessibility

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

9.10 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 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 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 Drives on RHEL4 Platforms

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

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

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

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

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

9.13 Supported Versions of IBM AIX Operating Systems

The AIX versions of the `pdscli` and `axiomcli` products distributed with the following Pillar Axiom software releases are supported only on AIX 5.3 TL5 and later versions of AIX:

Pillar Axiom 300	Pillar Axiom 500	Pillar Axiom 600
3.1 and later	3.1 and later	3.1 and later
2.0.3 and later 2.0.x	3.0.3 and later 3.0.x	-
1.8 and later 1.x	2.10 and later 2.x	-
1.7.3 and later 1.7.x	2.9.3 and later 2.9.x	-
1.6.10 and later 1.6.x	2.8.10 and later 2.8.x	-

9.14 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.

9.15 Resetting the Primary System Administrator Password

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

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

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

9.16 Uploading Software Update Packages over Slow Connections

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

9.17 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.

9.18 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.

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

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

9.19 WWN Designation Changed in Release 2.0

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

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

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

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

9.20 Some CLI Requests Have Been Removed

As of release 4.0, the following requests have been removed from the Command Line Interface (CLI):

- CreateLUNCopy
- ModifyLUNCopy
- DeleteLUNCopy
- GetAllLUNCopies
- GetLUNCopyDetails

- CreateVolumeBackup
- ModifyVolumeBackup
- DeleteVolumeBackup
- GetAllVolumeBackups
- GetVolumeBackupDetails

- CreateFileSystemCopy
- ModifyFileSystemCopy
- DeleteFileSystemCopy

- GetAllFileSystemCopies
- GetFileSystemCopyDetails
- PerformFileSystemCopyActivation
- PerformLUNCopyActivation

- PerformFileSystemBackupActivation
- PerformLUNBackupActivation
- RestoreFromVolumeBackup

- PerformBackgroundLUNCopy
- PerformBackgroundFileSystemCopy

- PerformTidyMediaPlacement
- SetValidate

As of release 4.1, the following request has been removed from the Command Line Interface (CLI):

- GetStorageConfigDetails (replaced by GetStorageConfig)

The above requests may appear to be supported in the `pdscli` or `supporttool` executable, but when these requests are invoked, the result will be the following:

> Error (4028):
>
> > The requested feature is no longer supported by the product.

9.21 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.

9.22 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.

9.23 Dynamic LUN Expansion Not Well Supported on Linux

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

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

9.24 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 typically causes 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.

9.25 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 CUs can be disabled in the GUI Global Network settings menu (System>Global Settings>Network).

Automatic failback of SAN Slammer CUs is always enabled.

9.26 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 Pillar the Pillar World Wide Customer Support Center for recovery assistance.

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

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

9.27 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.

9.28 Filesystem Fullness Impacts Write Operations

As a filesystem nears a “full” status, performance may slow down because finding storage for the data to be written takes more time.

9.29 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.

9.30 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.

9.31 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.

9.32 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 stops the resource and moves the volume to the new location. Attempting to use a logical volume while it is being moved can result in lost data.

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

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.

Important! If the LUN is a member of a SAN replication pair, you should isolate the pair before re-homing the LUN to a different Slammer.

9.33 SAN Replication Pair Isolation and Active Communication Channels

When isolating SAN replication pairs, the primary (source) and secondary (destination) Pillar Axiom systems must have an active communication channel between them. If the communication channel is lost between the primary and secondary systems, the replication pair will not be properly isolated at the secondary side.

In this situation should occur, restore the communications channel between the Pillar Axiom systems and re-issue the command to isolate the replication pair from the primary system.

9.34 Isolate SAN Replication Pairs Before Shutting Down

Pillar recommends that SAN replication pairs should be isolated before powering down a Pillar Axiom system. Doing so will reduce the communication between Pillar Axiom systems hosting the replication pairs during startup.

If SAN replication pairs are not isolated before a shutdown or a power loss does not allow time to isolate the pairs, the following procedure should be followed to shorten the time to start up the Pillar Axiom systems that have configured replication pairs:

1. Start up one Pillar Axiom system and then run the `start_session` command against that system.
2. When that command completes successfully, repeat Step 1 for all remaining Pillar Axiom systems participating in SAN replication.
3. After all systems have successfully started, run the `start_replication_pair` command for all replication pairs.

9.35 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.

9.36 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.

CAUTION! Running Slammer diagnostics, which tests and possibly modifies the contents of the battery-backed memory, can cause data loss if there is any data that has not been flushed to disk.

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.

Important! When running diagnostics on a SAN Slammer, the Pillar Axiom system assigns the Slammer ports "diagnostic" World Wide Names. When the diagnostics is complete, the administrator must reset the port names to their normal values.

9.37 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.

9.38 GUI Accurately Displays the Status of Components

In earlier releases, during system or core restart, all components would show a status of “Booting”. The GUI now more accurately displays the status of the components as they are discovered and initialized by the management software.

- During a Pilot restart:
    - The typical initial display shows the active Pilot CU as Booting. The standby Pilot CU may show Warning, Offline, or Booting, then transition to Online when start-up finishes.
    - Slammers initially show as Unknown as the Pilot software checks their functional status, then show Normal if the Slammers are still running.
    - Bricks show as Offline or Unknown and then transition to their true status as soon as the Slammer check completes.
- During a system restart:
    - 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 Booting, then Offline, and finally Online as the storage enclosures are discovered and initialized.

9.39 Cause of Degraded Status for a Volume Has Changed

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

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

9.40 A Drive May Display Blank Data in the GUI or Create an Administrator Action

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

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

9.41 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 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.

9.42 Alternative to Guided Maintenance Identify Step for Pilots

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

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

9.43 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.

9.44 Reuse of Pilot Control Units Can Cause Problems

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

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

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

9.45 Pilot Replacement Can Prevent SAN Replication From Restarting

SAN Replication persisted data must be made available on Pilot control unit (CU) replacements. If either or both Pilot control units must be replaced on a SAN Replication system, contact the Pillar World Wide Support Center for assistance.

9.46 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.

9.47 Replacing a Drive

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

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

  Important! You should contact the Pillar World Wide Customer Support Center before pulling a drive for test purposes.
- If a drive fails to be accepted into a Brick and the drive is set to Rejected status, do not attempt to use that drive. Contact Pillar Data Systems for another drive and for assistance.
- If an Administrator Action asking you to accept the 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 drive again. Contact Pillar Data Systems for another drive.

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

9.48 Moving Drives

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

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

9.49 Reseat Drives before Powering On New or Replacement Bricks

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

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

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

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

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

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

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

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

9.50 Testing Multiple Drive Failures

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

9.51 ACT LEDs on Drives Can Blink When Inactive

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

9.52 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.

9.53 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 will see the message: Topology Discovery Task.

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

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

These messages are normal and to be expected. 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.

9.54 Priority QoS Bands and Fibre Channel Bricks

As of release 3.1, Fibre Channel (FC) Bricks support the bleed up of lower-performance bands from SATA to FC storage in a mixed SATA/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.

As of release 4.1, such bleed up is no longer supported. Release 4.1 introduces the Storage Class feature, which only allows legacy volumes to grow or in-fill using the Brick type on which the volume was originally created. This feature does not automatically move legacy volumes to return them to their original Storage Class. If the Pillar Axiom system has a volume that has bled upwards, if you want to resolve it, contact the Pillar World Wide Customer Support Center.

CAUTION! After updating to release 4.1, a legacy thinly provisioned volume that has bled upward will not grow. This situation may cause data loss or pinned data if not resolved. This condition can be determined with a pre-upgrade audit and the Customer Support Center can offer assistance with resolving this type of issue.

All *new* volumes created in release 4.1 consume capacity only from within the Storage Class (Brick type) in which it was created.

9.55 Testing RAID Rebuild and Simulated Drive Fault

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

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

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

9.56 Brick Issues Can Cause a Slammer to Warmstart

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

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

9.57 Testing Brick Power Loss

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

Use the GUI Guided Maintenance screens for testing any power cycle of a solid-state drive (SSD) Brick.

CAUTION! Power cycling a Brick might cause the Pillar Axiom system to shut down. Before testing the power cycling of a Brick, contact the Pillar World Wide Customer Support Center to help you determine whether it is safe to power cycle the Brick. If the system configuration is located on the Brick being power cycled, the system will shut down.

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

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

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

CAUTION! As a safety measure, contact the Pillar World Wide Customer Support Center before accepting the Administrator Action. Because the Brick being recovered may contain persistence data (system configuration), accepting the Administrator Action may render the Pillar Axiom system inoperable, and cause it to shut down. (See also the **Important** notice in Section 9.60.)

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

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

9.59 Creating Logical Volumes Immediately After Replacing a Failed Drive

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

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

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

9.60 Persistence Vulnerability in Single-SATA Brick Systems

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

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

Tip: For assistance in determining which Bricks contain the Persistence store, contact the Pillar World Wide Customer Support Center.

9.61 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.

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

If the preceding operation succeeds, the Slammer CU status becomes Normal and is brought online automatically for SAN Slammers. NAS Slammer CUs are brought online automatically only when the global setting for failback of NAS resources is configured (see Enable Automatic Failback of NAS Control Units in the GUI).

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

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

9.62 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.

9.63 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.

9.64 Mapping UIDs and GIDs to Those of the NFS "anonymous" User

To map all UIDs and GIDs to the anonymous user, create the following special export under any filesystem on the desired File Server: `/PleaseEnableAllSquashFeature` (this special path does not have to exist).

Then, set the Anonymous User ID to the desired value.

Note: This capability functions in a way similar to the all_squash capability on Linux and similar NFS implementations.

The userID and GroupID for all users, including root, on any host defined with this export will be treated as though they were the configured anonymous user ID.

Normal user and root semantics and permissions will apply to any host that is not included in this export.

All exports listed after the `/PleaseEnableAllSquashFeature` export will map incoming UIDs and GIDs to the one specified under the UID field for the special export.

Setting the "Root Access" for a host will disable all mapping for that host, even though the special export exists.

9.65 Omitting UNIX Permissions When Copying Objects to a CIFS Share

When using NFS to copy files from a UNIX host to a CIFS share, to remove the UNIX permissions from the copied files, you have two choices:

- Use the proper options of the copy utility that explicitly does not preserve UNIX permissions during copy operations. For example, if you are using `rsync`, do not use the options `-o(owner)`, `-g(group)` or `-p(permissions)` for the copy. An example option is "`rsync -ulvt`".
- If you use a copy utility that does not have the flexibility of not preserving the UNIX permissions, you can modify the Unix permissions of the folders or files after the copy operation in the following way using the CIFS protocol:
    - From a Windows client, using an account with appropriate privileges, right click on the target parent folder (to where the UNIX folders and files were copied) and select **Properties** > **Security** > **Advanced**.
    - Check the following option: "Inherit from parent the permission entries that apply to child objects…."
    - Click **OK**.

Both methods overwrite the UNIX permissions of the copied files and folders with the Windows permissions of the parent folder.

9.66 Enabling the ABE and BTC Features for CIFS Connections

The two CIFS security features, access based enumeration (ABE) and bypass traverse checking (BTC), can be enabled through the CLI or the GUI by configuring the File Server supporting the desired CIFS connection.

- The ABE feature improves security by not displaying filesystem objects for which the client has no access rights. An ACL controls what a client can do, whereas ABE controls what the client can see.

  Note: ABE will reduce performance somewhat (because the system has to perform more security checks).
- The BTC feature relaxes filesystem security by removing access rights checking on the directories in the navigation path leading to the object that the CIFS client want to access. Access rights are still checked for the target object itself.

  Note: BTC will improve performance (because the system has to perform fewer security checks).

  BTC is available only in Windows 2003 R2 environments,

You can enable these features through the GUI by following these instructions:

1. In the top navigation pane, click **Storage**.
2. In the left navigation pane, click **File Servers**.
3. In the content pane, select the File Server for which you want to enable these features and click **Actions** > **Modify File Server**.

4. Click the CIFS tab.
5. In the **Server Comment** field, enable the feature.

   Choose one of:

   - To enable ABE, enter the following string: `PDS_ABE=1`.
   - To enable BTC, enter the following string: `PDS_BTC=1`.

   Note: To enable both features at the same time, enter both strings separated by a space.
6. Click **OK**.
7. Restart the CIFS File Server.

   Important! All current CIFS connections to this CIFS server will be lost when the server is restarted.

   To restart the CIFS File Server, perform one of the following actions:

   - Perform a significant configuration change to the CIFS File Server.
     Only that CIFS File Server is restarted. Often, the easiest configuration change is to change the WINS configuration of the CIFS server. Typical items that can be reconfigured without major impact to users are the list of WINS servers and the CIFS character set.

     Note: Changing the configuration of one CIFS File Server does not restart the other CIFS servers.
   - Force a warmstart of the Slammer control unit (CU).
     All File Servers on this CU are restarted.

     Tip: Pillar World Wide Customer Support Center can provide a Windows based utility for forcing a warmstart.
   - Perform a software update.
     All File Servers are restarted.

Pillar recommends the following:

- Enable the BTC feature to improve performance.
- Enable the ABE feature only when the security benefits outweigh any performance issues.

9.67 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.

9.68 Uploading Empty Files Through the GUI Not Allowed

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

9.69 Filesystem Quarantines

The filesystem quarantine feature can detect small inconsistencies in a filesystem. If an inconsistency is detected, this feature quarantines the affected resources rather than declaring the entire filesystem to be inconsistent and taking it offline.

After the resources are quarantined, the data is retained for read operations.

When the system quarantines a filesystem, you should contact Pillar World Wide Customer Support for help in restoring the filesystem.

9.70 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. These snapshots are recommended, 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. In addition, Pillar recommends that you create daily, weekly, and monthly snapshots.

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.

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 significant 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, it is recommended 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.

Important! If you perform a filesystem check on a filesystem, please contact the Pillar World Wide Customer Support Center to clear your logs once the situation has been resolved. Failure to clear logs may result in the Pillar Axiom system not being able to collect Slammer logs for future issues.

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.

9.71 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.

9.72 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.

9.73 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*.

A File Server may also be used in Windows 2000, Windows 2003, Windows 2008, and Windows 2008 R2 domains.

Furthermore, File Servers also support Windows 7 and Vista clients.

Additional CIFS features will be supported in future releases.

9.74 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.
- Refrain from applying domain local ACLs on share objects.
- Upgrade the domain from mixed mode to native mode.

9.75 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>$`. For example, to create a hidden administrative share for a filesystem named `foo`, create a share named `foo$`.

9.76 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.

9.77 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.

9.78 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.

9.79 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.

9.80 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 Failback fails, the system repeats the failover-failback sequence until the Slammer CU failure threshold (*see Section 9.81 below*) is reached. If that threshold is reached, the CU will be Disabled.
- 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.

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

9.81 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.

9.82 Support for the Pillar Axiom 10 GbE Feature

A correctly configured 10 GbE system should run at 10 Gb/s data rates.

9.82.1 Hardware Configuration

A variety of SFP and SFP+ modules exist in the field (1 GbE, 4 Gb Fibre Channel, 8 Gb Fibre Channel, and 10 GbE). Care must be taken because all SFP and SFP+ modules are physically interchangeable. For example, a 1 GbE module may be mistakenly inserted into a 10 GbE port, which may or may not operate. If that port does operate, it will run only at 1 Gb/s speeds.

Certain HBAs may require their own branded direct attach (copper) cables. Follow the manufactures installation guidelines.

For the Pillar Axiom system, optical testing was performed with an Intel X520 10 GbE NIM using Intel branded 10 GbE optical short range SFPs.

For the 10 GbE switch, optical testing was performed with a Cisco NEXUS 5010 10 GbE switch using Cisco branded 10 GbE optical short range SFPs. This switch was running the NX-OS 5.0(2)N1(1) operating system.

Note: Pillar does not recommend using a 1 GbE uplink on this switch.

Important! Pillar highly recommends that you only use HBAs and SFPs that are mentioned above.

Table 12 Supported 10 GbE optical cables (short range multimode fiber with LC-LC duplex connectors)

Item	Multimode fiber (MMF) optical cable, core diameter (micrometers)	Minimum modal bandwidth with short range 850nm laser (Mhz*km)	Operating range (minimum to maximum, in meters)	OM type	Fiber cable jacket color	Fiber cable core diameter markings
1	62.5	160	2 to 26	OM1	Orange	62.5/125
2	62.5	200	2 to 33	OM1	Orange	62.5/125
3	50	400	2 to 66	OM2	Orange	50/125
4	50	500	2 to 82	OM2	Orange	50/125
5	50	2000	2 to 300	OM3	Aqua	50/125

Note: The 82 meter OM2 optical cable was tested with an 80 meter cable spliced with a 2 meter cable.

Table 13 Supported 10 GbE direct attach copper cables

Item	Direct attach (copper) type	Operating range (maximum)
1	Passive	5 m

Note: *10 GbE network interface modules support only passive copper cables. Active copper cables are not supported.*

Important! For successful 10 GbE operation, only use respective cables that adhere to the maximum lengths in the above tables. If you exceed these maximum lengths, correct 10 GbE operation cannot be guaranteed.

9.82.2 Client Software Configuration

The 10 GbE feature has been tested using the following clients and drivers:

Table 14 Client operating systems and drivers

Operating system	Notations
Solaris 10 x86-64	ixgbe NIC driver 11.10.0 Rev=2005.01.21.16.34
Red Hat Linux U5	Drivers: o 2.0.75.1 - NAPI o 2.0.44 - k2
Windows 2008	Edition: R2 Enterprise x64

9.82.3 Diagnostic Commands

The `ifconfig` Slammer command (which can be entered on the **Support** > **Resolve Connectivity Trouble** screen) reports the media as Ethernet 1000baseTX full-duplex, which is incorrect.

To confirm the actual data rate at which a 10 GbE port is running, you can check its speed by navigating to the **Storage** > **Slammers** > **View Slammer Status** > **I/O Port Details** screen and scrolling down to the 10 GbE NIM. The port speed should show as 10 Gb/sec.

The same information can be retrieved by using the `GetSlammerStatistics` PDSCLI command.

9.83 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 (Disabled).

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

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

9.84 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.

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

9.85 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 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.

9.86 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.

9.87 Preventing Filesystem Corruption Due To Multiple Component Failures

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

9.88 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.

9.89 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.

9.90 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.

9.91 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.

9.92 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, some DMA vendors have 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.

9.93 NetBackup 6.5 Can Fail When Running on 64-bit Windows Servers

Some NetBackup 6.5 installations on 64-bit Windows Servers may experience failures during NDMP backup operations. The failures will typically occur at the very end of the backup job. To avoid these failures, upgrade to NetBackup release 6.5.4.

9.94 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.

9.95 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.

Note: Even though you can specify up to four ports to be aggregated, Pillar recommends as a best practice that no more than two ports be aggregated.

9.96 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.

Retina and other security scanners may incorrectly report Common Vulnerability and Exposures CVE-1999-0978, CVE-2000-0208 HTsearch CGI Script Detected on the management interface of the Pilot management controller. The resolution for those CVEs, as indicated in the vulnerability report, is Version 3.1.5 and higher of htdig. The Pillar Axiom system implements Version 3.1.6, which resolves this vulnerability and does not run htdig itself on the system but does run htsearch strictly for the help documentation. The htsearch is restricted to the help documentation and cannot access other areas of the Pillar Axiom system.

9.97 OpenSSH Has Some Vulnerabilities

The Pilot 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 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. The Support Administrator must issue a command to enable 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.

9.98 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.

9.99 Login to Oracle EM Plug-In Fails

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

9.100 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.

9.100.1 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.

9.100.2 Behavior

9.100.2.1 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.

9.100.2.2 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.

9.100.2.3 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.

9.100.2.3.1 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.

9.100.2.3.2 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.

9.100.2.4 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.

9.100.2.5 Notes

9.100.2.5.1 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.

9.100.2.5.2 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.

9.100.2.5.3 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.

9.101 Host Mappings Can Disappear Under Certain Circumstances

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

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

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

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

9.102 Overloaded Systems Can Result in Multiple System Warmstarts

The Pillar Axiom Pilot software is multi-threaded, and the number of threads available for handling system events and user-requested operations is limited. When a large amount of system activity requires Pilot intervention, this maximum number of threads can become fully utilized, leaving no threads available for user-requested operations. An excessive number of requests for action can result in a multiple failed tasks and multiple warmstarts. Should this situation arise, wait until the number of system events requiring Pilot action is reduced, freeing up the Pilot for user-requested operations.

9.103 IPStor Fails To Recognize LUNs After a Full System Upgrade

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

Pillar Axiom 300 and Pillar Axiom 500 systems can be upgraded to the Pillar Axiom 600 platform. However, as you do so, Bytes 16-31 of the Inquiry Data response (the "ASCII Product Identification" field) for each LUN will change from "Axiom 300" or "Axiom 500" to "Axiom 600". This change may cause some SAN appliances, such as IPStor, and possibly some hosts to not recognize the LUNs unless the appliance or host is reconfigured.

9.104 Getting Accurate SNMP Filesystem Performance Metrics

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

10 Technical Documentation Errata

The following sections describe topics in the technical documentation that were not able to be corrected in time for the current release of the Pillar Axiom 300, 500, and 600 system software.

10.1 *AxiomONE Administrator's Guide*

- The description of the RAID configuration for sequential writes needs to be changed in the following tables:
    - Table 20 (page 265)
    - Table 22 (pages 312 and 313)
    - Table 23 (page 334)

  In the above tables, replace the content in the "RAID Configuration in the Brick" column with the following:

  RAID 5

  Writes data in write-back mode to physical storage in full-stripe extents.

- On the following pages, the discussion of write caching needs to be changed:
    - Page 266
    - Page 314
    - Page 335

  For the above pages, replace the entire paragraph that starts with "The system stores all writes…" with the following (second sentence in the original is deleted):

  The system stores all writes of user data and system metadata in mirrored copies of the journal.

- On page 125, place the following paragraph immediately after the paragraph that starts with "NDMP-based backup and restore operations…":

  The DMA performs these backup and restore operations through the network interface modules (NIMs) in the Slammer storage controllers. These operations can utilize a locally attached tape device (if the NIM contains an HBA specifically for this purpose) or a tape device configured elsewhere in the customer's environment.

- On page 113, under the section titled "Validate File Integrity on a SecureWORMfs," replace the second paragraph with the following:

  Use this option when you want to verify the files outside of the weekly schedule or when you want to verify files on a specific directory (rather than the entire filesystem).

- On page 273, under the field named “Protected File Integrity Scan,” replace the first paragraph with the following:

  This scan checks all protected files on the filesystem and validates their data integrity.

- On page 96, after the first paragraph, add the following notice:

  Note: When you change the Storage Class of a logical volume, the volume is migrated to the new Storage Class but any existing clones of that volume are not.

- On page 106, before the first step, add the following notice:

  Note: When you change the Storage Class of a logical volume, the volume is migrated to the new Storage Class but any existing clones of that volume are not.

- On page 263, in the “Storage Class” section, add the following notice after the last bullet:

  Note: When you change the Storage Class of a logical volume, the volume is migrated to the new Storage Class but any existing clones of that volume are not.

- On page 311, in the “Storage Class” section, add the following notice after the last bullet:

  Note: When you change the Storage Class of a logical volume, the volume is migrated to the new Storage Class but any existing clones of that volume are not.

10.2 *AxiomONE Replication User Guide and Reference for SAN*

- On page 14, replace the entire content under “Feature Licensing” with the following:

  The AxiomONE Replication for SAN utility requires various licenses to be installed on the Pillar Axiom systems that participate in replication operations.

  - *AxiomONE Replication for SAN. This license must be installed on all Pillar Axiom systems designated as the primary Pillar Axiom system for a replication pair.*
  - *Fibre Channel or iSCSI Protocol. A SAN Slammer ships with the appropriate protocol, or both, depending on your network requirements. A license for at least one of these protocols is required to support path configuration and other replication operations.*

- On page 49, add the following note after the last paragraph:

  ***Note:** At minimum, each SAN Slammer control unit (CU) on the primary Pillar Axiom system must be connected by at least one path to each SAN Slammer CU on the secondary Pillar Axiom system before the replication pair can be successfully started. Thus, a minimum of four paths is required in a channel linking two single-Slammer Pillar Axiom systems.*

10.3 *Pillar AxiomONE CLI Reference Guide*

On page 38, change the syntax for the `clonefs -add` command to the following:

```
clonefs –add –sourcefileserver source-file-server
  –sourcefilesystem source-filesystem
```

10.4 Pillar Axiom NDMP Integration Guide for NAS Systems

On page 7, place the following paragraph immediately after the paragraph that starts with "NDMP-based backup and restore operations…":

> The data management application (DMA) performs these backup and restore operations through the network interface modules (NIMs) in the Slammer storage controllers. These operations can utilize a locally attached tape device (if the NIM contains an HBA specifically for this purpose) or a tape device configured elsewhere in the customer's environment.

10.5 *Pillar Axiom CIFS and NFS Multi-Protocol Planning Guide*

At the bottom of page 64, add the following content:

> *When a CIFS client creates a directory having permission inheritance, any NFS operation to create a new object in that directory will not control the initial permissions of the new object. Instead, the initial permissions of the new object are determined by the inheritable Access Control Elements (ACEs). An inheritable ACE applies to more than just the current folder. For example, if an ACE applies to "This folder and any subfolders," it is an inheritable ACE.*
>
> *The following list summarizes the rules for both NFS and CIFS:*
>
> - ***New object created using NFS***
> - *If the existing directory does not have inheritance, the new object starts with NFS permissions reflecting the user's UMASK.*
> - *If the existing directory does have inheritance, the new object starts with security that is inherited from the existing directory.*
> - ***New object created using ordinary CIFS operation***
> - *If the existing directory does not have inheritance, the new object starts with default CIFS permissions.*
> - *If the existing directory does have inheritance, the new object starts with security that is inherited from the existing directory.*
> - ***New object created using unusual CIFS operation that specifies the security for the new object***
>
> *The new object starts with the security provided by the operation, regardless whether the existing directory has inheritance.*

10.6 *Pillar Axiom 300 Advanced Hardware Installation Guide*

On page 55, add the following notice:

CAUTION! *Each control unit (CU) in a Pillar Axiom 300 Slammer has only one power supply. When half the power to a Pillar Axiom 300 system is lost, improper routing of the power cables from the Slammer and Pilot to the AC power supply can cause the private management interface (PMI) Ethernet to hang. To avoid this situation, route the power cables as follows:*

- *Pilot CU 0 and Slammer CU 0 connect to the same AC power supply (power supply A).*
- *Pilot CU 1 and Slammer CU 1 connect to the same AC power supply (power supply B).*

10.7 *Pillar Axiom 600 Hardware Installation Guide*

- On page 175, replace the column headings of Table 52 with the following headings:

Multimode Fiber (MMF) optical cable type core diameter (micrometers)	Minimal modal bandwidth with short range 850nm laser (MHz*km)	OM type	Maximum cable length

- On page 175, add the following notice after Table 52:

  Note: *The optical cable is short range multimode fiber with LC-LC duplex connectors.*

- On page 175, add the following notice after Table 53:

  Note: *10 GbE network interface modules support only passive copper cables. Active copper cables are not supported.*

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

On page 31, Step 5 should read as follows:

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

10.9 *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`.

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

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

- “Configure the iSCSI Software Initiator”
- “Start the iSCSI Software Initiator Service”

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

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