| Pillar Axiom® 500 | Document Number: 4420-00026-2200 |
|---|---|
| | Document Title: Customer Release Notes |

**Revision History**

| Rev Description | Rev Date | Effective Date |
|---|---|---|
| Release 03.00.00 | 2008-03-12 | 2008-03-14 |
| Release 02.08.00 | 2008-01-10 | 2008-01-14 |
| Release 02.07.00 | 2007-10-29 | 2007-11-05 [expand list of fixed issues] |
| Release 02.07.00 | 2007-10-01 | 2007-10-22 |
| Release 02.06.00 GA | 2007-06-29 | 2007-06-29 |
| Release 02.06.00 Beta | 2007-06-08 | 2007-06-08 |
| Release 02.05.00 | 2007-04-13 | 2007-04-13 |
| Release 02.04.00 | 2007-02-11 | 2007-02-12 |
| Release 02.03.00 | 2006-11-17 | 2006-11-20 |
| Release 02.02.00 | 2006-10-31 | 2006-11-4 |
| Release 02.01.00 | 2006-10-11 | 2006-10-11 [corrected Slammer PROM version] |
| Release 02.01.00 | 2006-09-28 | 2006-09-29 |
| Release 02.00.00 | 2006-08-17 | 2006-08-17 |

# 1 Purpose

This document describes new features, capacities, configuration requirements, operating constraints, known issues and their workarounds, and other items for Release 03.00.00 of the Pillar Axiom 500 storage system. The document covers hardware, firmware, software, cabling, and documentation. The information provided is accurate at the time of printing. Newer information may be available from your Pillar Data Systems authorized representative.

# 2 Product Release Information

Release 03.00.00 is a major release of the unified NAS/SAN Pillar Axiom storage system.

## 2.1 System Enhancements

This update provides substantial quality improvements and feature enhancements to the Pillar Axiom 500 storage system. Release 03.00.00 provides improved system efficiency and robustness and optimized Quality of Service (QoS), as well as including a rollup of all defects fixed in customer patches up to Release 02.08.00.

- The new Pooled RAID 10 feature provides QoS optimization and performance enhancements for random-write intensive applications. This feature uses parallel mirrored-write operations.
- The new Thin Provisioning feature allows a logical volume (LUN, filesystem, and their clones) to be created that appears to be much larger than the physical storage actually backing it. As additional capacity is needed for the volume, the system dynamically allocates additional storage to it.

  **Note**: Customers who have existing NAS systems who want to update their system to Release 03.00.00 should contact their 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.
- Clones of filesystems and LUNs are now writeable and use space-efficient, partial-block snapshot technology.
- In many cases, the system will now quarantine inconsistent blocks in a filesystem, allowing the filesystem to remain online. Administrators can take a number of actions, including for example the deletion of the quarantined data or the scheduling an FSCK (check-only mode or full mode).
- The restriction in Pillar Axiom NAS systems that require all File Servers except the first to have VLAN tags has been removed.
- SecureWORMfs filesystems now include the ability to audit all content transactions and have an immediate verification option. Also, Standard SecureWORMfs filesystem can be changed to a Compliance variant where the journals are updated on disk.
- The SecureWORMfs feature allows the Standard variant to use battery-backed memory journaling, which speeds up the initial data population.
- Improvements to the Call-Home feature include a repository for storing multiple Call-Home and log collection bundles.
- Administrators can release stale NFS NLM locks (which a client may have initiated but failed to release) on any and all files for a given File Server.
- This release greatly increases the number of hard links that can point to original source files.
- This release enhances CIFS security by allowing administrators to create a list of preferred LDAP and Kerberos servers or Domain Controllers for authentication.
- File Servers can now share VLANs, allowing File Servers to be on the same VLAN with or without tags.

- This release provides full support for these additional locales:
    - Europe: CIFS Windows locales (ISO8859)
    - Japan: CIFS Windows locales (Shift JIS) and NFS character sets (EUC-JP)
- VSS speed has improved to synchronize 20 snapshots in less than 10 sec.
- Three and four Slammer systems are now supported.
- 4-port GbE network interface modules (NIMs) are now available. These NIMs are available with either copper or optical GBICs.

  **Note**: 4-port GbE NIMs are supported only for Pillar Axiom systems running Release 03.00.00 software. For this release, if a 4-port GbE NIM is used in a system, all NIMs in the system must be 4-port GbE.

  **Tip!** Link aggregation is supported for 4-port NIMs in the following port configuration sets:
    - [0,1]
    - [2,3]
    - [0,1,2]
    - [0,1,2,3]

Note: Release 03.00.00 drops the Triple Redundancy QoS option for *new* volumes. However, using Double Redundancy in a Pooled RAID 10 array effectively provides quadruple redundancy, if that is desired.

For the list of defects that this release resolves, see Table 6 beginning on page 26.

## 2.2 Available Licenses

Licenses available for this release include the following:

- CIFS
- NFS
- SNMP
- NDMP
- Pooled RAID 10
- SecureWORMfs
- Snap FS
- Clone FS
- Clone LUN *(formerly Snap LUN)*
- Thin Provisioning: FS
- Thin Provisioning: LUN
- Volume Copy
- Volume Backup/Restore
- Volume Copy and Backup

- Fibre Channel Protocol
- iSCSI Protocol
- Support Tool
- Axiom File Replicator
- Axiom Volume Replicator
- VTL Appliance

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

## 2.3 Pillar Axiom Software Update

This release may be installed through the GUI Software Update process. Be sure to select all of the software components that have selection boxes available; otherwise, the update may not proceed due to a potential incompatible combination of components.

**Important!** After updating the system software to Release 03.00.00, you cannot reverse the update and downgrade to an earlier release level.

**Important!** Before you begin the software update process for Release 03.00.00:

- *Ensure no background processes are running.*
  You can check for running background processes from the system GUI by pressing the "Display All Background Processes Running" button in the lower right corner of each GUI screen.
- *Do not initiate any system or storage actions of any kind.*
- Perform the update during a maintenance window during which all user applications are stopped. Updating a Pillar Axiom system to Release 03.00.00 disrupts all data paths.

Also, please be aware of the following notes before you begin the software update process:

- For NAS/SAN or SAN systems at 01.07.xx, the system cannot be updated directly to 03.00.00. A system must be running 02.07.00 (or later) before it can be updated to 03.00.00. Contact Pillar Technical Support for assistance in updating any system below release 02.03.00. Do not attempt to perform this update without this assistance.
- **Important!** If your system contains Fibre Channel Bricks and is running Pillar Axiom software below release 02.04.02, contact Technical Support to update the system.
- **Important!** When you begin the update process, you must select *all* software modules.

After a successful update, all system components will report Normal in the GUI and user data will be available. Users can start up their applications to access this data.

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

### 2.3.1 Initialization of Unused Capacity

Immediately after the update, all *unused* capacity will be temporarily unavailable. Available unused capacity will begin to grow immediately as the system initializes (zeroes) the unused space.

**Tip!** The GUI displays the available space dynamically while zeroing proceeds. Refreshing the Storage > Usage Summary screen shows the free space in the system as it becomes available.

While the system is initializing the unused capacity, Pillar recommends that storage administrators avoid performing the following actions:

- Create new logical volumes (filesystems and LUNs).
- Clone existing volumes.
- Back up or copy existing volumes.

Similarly, users should avoid uploading large amounts of data.

The initialization process occurs in parallel on all Bricks in the system. As an example, a system with 1TB drives and no load would take approximately 5 hours to initialize.

- The amount of time to initialize the unutilized space is load dependant. If an application load is applied to the system, initialization can take up to 8 times longer than an idle system.
- The more Bricks and the smaller the disk drives the faster initialization will be.

### 2.3.2 Change in the Amount of Reported Free Capacity

After zeroing completes, the reported free capacity will be larger than what was reported prior to updating to Release 03.00.00. The free capacity reported in Release 03.00.00 is raw capacity and does not account for parity space used by RAID 5 arrays, hence the larger number.

Prior to this release, the reported free capacity was the amount of space available for storing user data and had already subtracted RAID 5 parity space. Because this release supports Pooled RAID 10 in addition to RAID 5, raw capacity is now reported.

### 2.3.3 Software Versions for this Release

Software provided in this release includes the versions listed below. After updating your system, check your software versions in the GUI by going to Support > Software Modules.

| | |
|---|---|
| Pilot OS | 03.00.00 |
| Pilot Software | 03.00.00 |
| Slammer PROM | 03.00.00 |
| Slammer Software | 03.00.00 |
| SATA Brick Firmware | 07.00.45 |
| Fibre Channel Brick Firmware | 00.60.45 |
| Brick Disk Drive Firmware | *(Version depends on the disk drive capacity and whether it is the spare or array disk drive.)* |

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

**Important!** Starting with the 2.x releases, the WWN was changed to use a common base World Wide Node Name. When updating a Pillar Axiom system from 1.7.x to 2.x, be prepared for a change in how WWNs are managed. For more information, see Section 9.14 on page 43.

### 2.4 AxiomONE Path Manager Software

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

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

The current release of APM for SAN hosts varies by platform, as shown in Table 1.

**Table 1 APM support**

| Platform | APM release | Notes |
|---|---|---|
| AIX | 2.1.4 | Supports SAN boot on AIX 5.2 and 5.3. |
| HP-UX | 2.0.3 | - |
| LINUX | 2.1.8 & 2.2.2 | Red Hat Enterprise Linux 3 |
| | 2.1.3 & 2.2.2 | Red Hat Enterprise Linux 4<br>Novell SUSE Linux Enterprise Server 9 |
| Solaris | 2.0.2 | Solaris 10 |
| | 01.04.05 | Solaris 8 |
| | 2.0.2 | Solaris 9 |
| Windows | 2.3 | - |

Note: Unless otherwise explicitly stated in your APM Release Notes, there are no co-requisite relationships between the AxiomONE Path Manager and Pillar Axiom 500 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.

## 3 Terms and Conditions of Use

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

# 4 Support

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

**Table 2 Contact information**

| For help with… | Contact… |
|---|---|
| Technical Support | U.S. and Canada: 877-4PILLAR (877-474-5527)<br>Europe: +800 PILLAR FS (+800 74 55 27 37)<br>Asia Pacific: +1-408-518-4515<br>South Africa: +0 800 980 400<br>Have your system serial number ready.<br>Email: support@pillardata.com<br>Web: support.pillardata.com |
| • Implementation assistance<br>• System information<br>• Enhancement requests | sales@pillardata.com.<br>USA: 1-877-4PILLAR (1-877-474-5527)—request Sales at the prompt.<br>International: +1 408 503 4200 |
| Documentation improvements and resources | docs@pillardata.com.<br>www.pillardata.com/techdocs—log in with your username and password. |

## 4.1 Supported Hardware Components in a Pillar Axiom System

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

## 4.2 Access to Pillar Axiom Systems

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

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

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

### 4.3 Download Software or Firmware Updates

To download software or firmware updates:

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

Note: To obtain a license to enable a feature that is in addition to those that you initially purchased, contact a Pillar sales representative. (See Table 2  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 information regarding the primary features of a Pillar Axiom storage system and how to configure them:

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

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

## 5 Pillar Axiom System Limits

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

**Important!** Use care when operating a system that has been configured to run at or near the system operating limits. The system may exhibit anomalies when all limits are exercised concurrently. Also, the time to start up a Pillar Axiom 500 system 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 500 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 3 Operating limits of a Pillar Axiom 500 NAS system**

| Item | Description and Range |
|---|---|
| File Servers | Maximum = 8 per Slammer<br>Note: In multi-Slammer systems, a File Server's virtual interfaces (VIFs) can be configured on multiple Slammers. The presence of VIFs is what counts against the above limit. Such a File Server is considered to be present on each Slammer on which it has VIFs. |
| VIFs per File Server | • Minimum = 1<br>• Maximum = 32 |
| VIFs per Slammer port | Maximum = 16 (across all File Servers) |
| VLANs | No restrictions |
| Static and default network routes for each File Server | Minimum = 0<br>Maximum = 5 |
| NIS alternative file size | Up to 50 MB |
| Volume groups | From 1 to 5000, out to four levels |
| Filesystems | Minimum = 1<br>Maximum =<br>• 128 per Pillar Axiom 500 system<br>• 128 per NAS Slammer |
| Filesystem size | Minimum of 1 GB and 50% of maximum capacity<br>Minimum growth increment of 1 GB<br>Maximum = system capacity |
| Snapshots | Maximum =<br>• 250 per filesystem<br>• 10,000 per Pillar Axiom 500 system |
| Volume Backups | Maximum =<br>• 1024 per Pillar Axiom 500 system<br>• 1024 per filesystem<br>Maximum concurrent = 5<br>Maximum outstanding I/O requests = 128 across all backups |
| NDMP | Maximum concurrent backup/restore sessions per system = 5 |
| NFS exports | Maximum = 1000 per File Server |

| | |
|---|---|
| NFS host entries | Maximum = 4000 per File Server |
| CIFS shares | Maximum = 128 per File Server |
| User security groups | Maximum = 254 per CIFS user |
| CIFS connections | Maximum per Slammer:<br>• 1200 for 12 GB memory (total)<br>• 3000 for 24 GB memory (total) |

**Table 4 Operating limits of a Pillar Axiom 500 SAN system**

| Item | Description and Range |
|---|---|
| SAN LUNs | Maximum =<br>• 2048 visible per system<br>• 256 visible per host<br>• 1024 visible per SAN Slammer |
| SAN LUN size | Minimum of 1 GB and 50% of maximum capacity.<br>Minimum growth increment of 1 GB<br>Maximum = system capacity |
| Volume Copies<br>(full block snapshots) | Maximum =<br>• 128 per Pillar Axiom 500 system<br>• 12 active per LUN |
| Clone LUNs<br>(partial block snapshots) | Maximum = number of unallocated SAN LUNs, up to 1024 |
| iSCSI | Maximum =<br>• 256 TCP connections per iSCSI port<br>• 256 iSCSI Initiators per iSCSI port<br>• 32 persistent reservation registration keys per LUN<br>• 512 simultaneous commands per iSCSI port |

### 5.2 Slammer and Brick Configuration Limits

| Number of Slammers | Minimum number of Bricks‡ | | Maximum number of Bricks |
|---|---|---|---|
| | Supported | Recommended | |
| 1 | 2 | 3 | 32 |
| 2 | 4 | All NAS or all SAN Slammers = 5<br>NAS / SAN combo Slammers = 6 | 64 |
| 3 | 6 | 8 | 64 |
| 4 | 8 | 8 | 64 |

For Fibre Channel (FC) Bricks specifically:

- Fibre Channel (FC) Brick storage enclosures normally come in pairs of one FC RAID Brick plus one FC Expansion Brick. A FC RAID Brick however can be installed without an Expansion Brick.
- One-Slammer systems have a limit of 16 FC Bricks. Multi-Slammer systems can have up to 32 FC Bricks.
- Both SATA and FC Bricks may co-exist in the same Brick string subject to current configuration recommendations.
- A given Brick string can contain up to a total of 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 and SATA Bricks, see Appendix C in the *Pillar Axiom 500 SSF Cabling Reference*.

## 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.

‡ The minimum number of Bricks in any configurations is two of the same type.

When logging into the AxiomONE Storage Services Manager using Secure HTTP, you may see warnings that the server certificate is not issued by a trusted authority. The server certificate is installed and signed by Pillar Data Systems during the manufacturing process.

## 6.2 Network Requirements

### 6.2.1 Pilot Network Requirements

The Pilot management controller requires:

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

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

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

### 6.2.2 Slammer Network Requirements

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

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

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

## 6.3 NDMP Requirements

For a list of data management applications (DMAs) that Pillar Axiom systems support, see the *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 500 system is the Pilot management interface. Data movement is performed over the data path interfaces. Be sure that any external NDMP backup servers are able to reach the Pilot management IP addresses.

### 6.3.3 Virtual Interface (VIF)

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

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

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

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

### 6.3.4 Fibre Channel (FC) Tape Library LUNs

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

## 6.4 Power-Off Requirements

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

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

## 6.5 Power Cycling

Contact Pillar Technical Support before power cycling a Pillar Axiom 500 system except in the event of an emergency, in which case, drop all power and then contact Technical Support. Contact Technical Support 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 Pillar Technical Support.

# 7 Known Issues

The Pillar Axiom 500 server issues listed in Table 5 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 2 Contact information).

**Table 5 Known Pillar Axiom 500 server issues**

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| ALL | If a non-disruptive software update is performed, it is possible for the Bricks to end up in an unusable state with RAID firmware at two different levels. | Do not attempt to recover on your own. Instead, contact Pillar Technical Support for assistance.<br>This issue will be fixed in a future release. |
| | For systems containing a combination of SATA and FC Bricks, for all but Random Write QoS, one should be able to create a Premium-priority volume that not only uses all remaining FC space (if any) but also uses available SATA space as needed. However, attempting to create a Premium volume that initially exceeds the amount of remaining FC space results in an "insufficient capacity" error.<br>Note: This issue does not apply to volumes created with Random Write QoS, which can never span different Brick types. | When creating volumes with a Priority setting of Premium, do not attempt to create a volume that exceeds the remaining capacity of the Premium priority band.<br>Or, create the volume in non-Premium bands that have available capacity.<br>This issue will be fixed in a future release. |
| | For systems containing a combination of SATA and FC Bricks, the "Option one" table on the Optimizer page in the GUI suggests that you can create a non-Premium volume that consumes all remaining SATA space (if any) and, if needed, available FC space.<br>The documentation also says that this is possible and advises using an Archive setting in order to use as much capacity as possible, regardless of Brick type.<br>However, creating such a volume results in an "insufficient capacity" error.<br>Note: This issue does not apply to the "Option two" (Random Write) table on the Optimizer page, which shows the FC space and SATA space as distinct. Volumes created according to the terms of this table cannot span different Brick types. | When using the "Option one" table to create a volume with a Priority setting of Archive, Low, Medium, or High, perform the following steps to ensure that the new volume uses only SATA space:<br>1. Find the Optimizer value for the desired Redundancy and Priority.<br>2. Find the corresponding Premium value for the same Redundancy.<br>3. Subtract this value from the value in Step 1.<br>4. Make sure the initial capacity of your new volume is smaller than the result of Step 3.<br>This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | The pdscli command PerformRemoveBrickFromConfigTask doesn't work. This failure also precludes using Guided Maintenance to remove a Brick from the system. | Do not attempt Brick removal without first contacting Technical Support.<br>This issue will be fixed in a future release. |
| | Very large system information or Call-Home files cannot be downloaded to a client system using the web download function. You get an error popup:<br>A system error has occurred:<br>Cannot download file. Try again or contact Technical Support. | Try using the transfer function.<br>Alternatively, reduce the amount of information in the bundle by selecting "most recent" or by breaking the information into two bundles: one with all of the Bricks and another with all of the Slammers.<br>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:<br>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.<br>This issue will be fixed in a future release. |
| | In some cases, during power on of a Brick containing 1 TB disk drives, it is possible for a disk drive to lose connection to the Brick, resulting in the disk drive being marked Critical. | Replace the disk drive.<br>**Important!** If more than one disk drive fails in a Brick, do not attempt replacement; instead, contact Technical Support.<br>This issue will be fixed in a future release. |
| | The system does not properly keep track of objects with a High Throughput profile. As a result, a number of operations, when performed on filesystems or LUNs that have this profile, may not be done correctly. | High Throughput is intended for special applications only. If you intend to use this profile, contact Technical Support.<br>This issue will be fixed in a future release. |
| | If a disk drive is performing marginally, the system does not notify the storage administrator until the error-rate becomes high enough that the drive is taken offline. At that point, the spare disk drive is used automatically. | This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | The PDSCLI command "GetStorageConfigDetails" returns raw capacities based on a RAID 5 geometry. This does not account for capacity requirements when using Pooled RAID 10. The RAID 10 raw capacity is 1/2 of the listed capacities returned by this command. | When using the PDSCLI command "GetStorageConfigDetails" on a Pooled RAID 10 system, divide the listed capacities by 2 to get the actual capacity. .<br>This issue will be fixed in a future release. |
| | When updating Brick software while the Brick has a heavy I/O load, it is possible for the Brick to become unresponsive for a short time. | Let the process complete.<br>If possible, the update should be performed during a maintenance window when I/O is light.<br>This issue will be fixed in a future release. |
| | The Configuration Wizard does not detect whether the system is SAN-only or NAS-only. Therefore, the Wizard might seem to allow the user to create objects and attempt operations that are not appropriate for the actual system. Eventually any attempt to create actual objects on the system that the system doesn't support will end in failure. | Do not use the Configuration Wizard to create objects that the system does not support. Be aware that the Configuration Wizard will allow you to create unsupported objects.<br>This issue will be fixed in a future release. |
| | When an administrator requests an immediate software update, there is no warning that an already-scheduled software update might start up and interfere with the immediate update while it's in progress. | When requesting an immediate software update, be certain that a software update is not already scheduled to take place within the next hour or so.<br>This issue will be fixed in a future release. |
| | When changing the Pilot management IP from static to DHCP in the GUI, the change does not get committed. | Run the CLI command ModifyManagementConfig with DHCPEnabled set to true and the pilot1 and pilot2 IPAddress fields set with proper static values.<br>This issue will be fixed in a future release. |
| | “Revert Software to Previous Version” in the GUI fails to downgrade Pilot OS and Pilot Software. | Do not use “Revert Software to Previous Version”. Upgrades to release 03.00.00 is not reversible. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | When rescheduling a software update from the GUI, the active Pilot control unit (CU) will go down and fail over to the buddy CU. | If you need to reschedule the update, delete the old scheduled update and create a new one.<br>Alternatively, perform an immediate software update.<br>This issue will be resolved in a future release. |
| | If you try to replace a Brick chassis using Guided Maintenance, the Brick will go Critical with all components in an unknown state. You may also see an error “Required Brick tag not found”. | Replacing Brick and Slammer chassis are not supported in this release. If you have attempted to replace a Brick chassis, contact Technical Support for assistance in recovering.<br>This issue will be resolved in a future release. |
| | Under rare conditions after a software update, the option to revert to the previous version may not be displayed in the GUI. | Contact Technical Support for assistance if you need to revert the software to the previous version. (Not applicable to Release 03.00.00 because that release is one way.)<br>This issue will be resolved in a future release. |
| | If a software module update is selected that is already installed on the system, the Pillar Axiom system will still indicate that a software update is in progress for several minutes before returning to normal status. | This issue will be resolved in a future release.<br>Note: If you select RAID firmware, the update process will re-install the firmware (even if it hasn’t changed). |
| | On a system with a considerable number of outstanding Administrator Actions, the GUI screens will display very slowly. Attempting to display the active Administrator Actions may result in the display hanging. | Contact Pillar Technical Support for assistance in determining its cause and in removing the Administrator Actions.<br>This issue will be fixed in a future release. |
| NAS | If a CreateCloneFS request is followed by a second CreateCloneFS request on the same filesystem while the first one is still in progress, the second request usually fails (as designed), but the GUI does not display an error dialog on the failure. | When you need to create a Clone FS on a filesystem or a Clone FS, before trying to create another Clone FS on the same filesystem, wait until the status of the created Clone FS becomes Available. You can check the status using the GUI or the CLI request GetAllCloneFSs.<br>This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | 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.<br>The domain controller should not block any port for LDAP services.<br>This issue will be fixed in a future release. |
| | Non-disruptive software updates on 03.00.00 systems that have many NFS exports and hosts or CIFS shares can take much longer to complete than for example a system restart or disruptive update of the same system. | Contact Technical Support for help.<br>This issue will be fixed in a future release. |
| | A restore of a filesystem from a snapshot can appear to be complete to the user from the system user interfaces, but internally the system might still be restoring the filesystem. If the user initiates a second restore to the same filesystem while the system is still working on the first restore, the second request will hang indefinitely and always show up as a task in progress in the GUI. | Avoid restoring to the same filesystem within too short a time. (The time to restore a filesystem is a function of the size of the filesystem, CPU load on the Slammer, and other things.) If you submit a second request and it hangs, contact Technical Support for help in recovering.<br>This issue will be fixed in a future release. |
| | When using the local group assignment for the administrator group, the users in this group will get additional privileges equivalent to "root". | This issue will be fixed in a future release. |
| | In very rare conditions where the system previously took a filesystem offline due to internal issues that were eventually resolved, the system continues to report the filesystem as Offline, even though the filesystem is online and serving data. | Contact Technical Support for help in clearing the spurious Offline status.<br>This issue will be fixed in a future release. |
| | When a CIFS server has joined a domain in ADS mode and a user successfully authenticates using ADS and connects to that server, if the administrator then assigns CIFS local group and the user subsequently tries to make use of the privileges as part of CIFS local group assignment, such as trying to access a folder for which he or she did not previously have permissions, the attempt will fail. | Log out, log back in, and then retry the operation.<br>This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | If a single disk drive fails in a FC RAID or Expansion Brick, the failure may cause the entire Brick to go offline and display as Critical. | Contact Technical Support for help in recovering the Bricks.<br>This issue will be fixed in a future release. |
| | When backing up a directory, if the following is true, the backup operation will fail:<br>• The traverse bit permission is not set for the directory.<br>• The user running the backup belongs to the “local backup group”.<br>• The user has no permissions in that group. | Use the Administrator account or the Domain Backup user account to run the backup application.<br>This issue will be fixed in a future release. |
| | When running an NDMP file-based restore to a Compliance WORM filesystem, the data transfer may run abnormally slow. This is more likely to happen when restoring many small files (< 16 KB) where the ratio of directories to files is extremely high. | Restore to a Standard WORM filesystem or to a non-WORM filesystem. If that is not possible, select only the files needed for restore instead of running a full restore.<br>This issue will be fixed in a future release. |
| | The CIFS client attempts to open a file with delete access, but the Pillar Axiom File Server rejects the request, because the server does not allow files to be renamed or moved if they have the "Read-only" attribute set. This action is not consistent with the operation of a Windows server. | If possible, remove the "Read-only" attribute before renaming or moving files.<br>This issue will be fixed in a future release. |
| | After running the ebifsx test suite, removing all files and directories created by the test might fail with a message about a directory being "not empty". This directory contains hidden files, which an `ls` command will not list. | If you want to continue using the filesystem, create a new directory (named for example “lost-found”) and move the "bad" directory there.<br>Alternatively, if the filesystem is a normal one, you can delete the entire filesystem. However, if the filesystem is a WORM filesystem, you cannot delete it when the number of files in that filesystem is greater than zero.<br>This issue will be fixed in a future release. |
| | A File Server can experience out of service for up to 10 min with NFS when large writes to filesystems are accumulated. When this condition occurs, the system warmstarts to correct the condition. | This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | During heavy write I/Os to a source filesystem or Clone FS, the repository used to hold the changes may get full. In such cases, to keep the source filesystem online, the system breaks the association between the filesystem and its clones, and the clones will go Offline with pinned data. All further writes to the Clone FSs will fail. | These clones need to be deleted after discarding the pinned data.<br>This situation can be avoided by having a big enough repository. The respository size needs to be planned ahead during Clone FS creation because repository size cannot be changed later.<br>This issue will be fixed in a future release. |
| | Unchecking 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.<br>This issue will be fixed in a future release. |
| | If a snapshot is restored after the filesystem FSCK is performed, there is a chance the snapshot can still have the same corruption. | Contact Technical Support for help in restoring the filesystem.<br>This issue will be fixed in a future release. |
| | Cannot use NDMP to perform a block backup a Clone FS. | Block backups of clones are not supported. Instead, use Backup to Disk Now.<br>This issue will be fixed in a future release. |
| | Because of a network issue, the request from the Pillar Axiom system generated no response, eventually causing the system to fail a health check after the health check timed out. | Check the network and Domain Controller if this happens repeatedly.<br>This issue will be fixed in a future release. |
| | A write request may be rejected with a quota error if a write operation in its entirety exceeds the hard quota limit. | Avoid exceeding quota soft limits. When these limits are reached, events are logged warning of soft limits being exceeded.<br>This issue will be fixed in a future release. |
| | When CIFS is configured with Account Mapping enabled and with an NIS server, CIFS server attempts to use NIS to map CIFS user names to NFS user names. This mapping occurs during authentication processing while mapping a drive. If the NIS server does not respond promptly, the lower layers of code will set a timer and retry. In certain cases, the retry time will exceed the client’s timeout setting, and the client will drop the connection, failing the authentication request. | Ensure the NIS server is actively serving requests. If the NIS server is not reliable, consider removing it from the Pillar Axiom configuration and using local files for name resolution. If Account Mapping is not mandatory in the customer environment, disable it.<br>This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | A Pillar Axiom NAS system may not finish a software update successfully when it is under extremely heavy loads or has Brick issues that prevent I/O requests from completing in a reasonable amount of time. In such cases, the NAS control units may fail. | Stop all I/O traffic to the Pillar Axiom system before updating the system software.<br>Note: Refer to Section 2.3 for more information.<br>This issue will be fixed in a future release. |
| | When using CIFS to copy files to a Pillar Axiom filesystem, the target files will have an associated security descriptor created. This behavior occurs even when the source file did not have an explicit security descriptor. As a result, the space required on the target may be larger than the space required on the source filesystem. | Reserve additional space on the target filesystem when copying files using CIFS.<br>This issue will be fixed in a future release. |
| | When using a device with less than 1 Gbps link speed to connect to a NAS Slammer, the I/O Port Details screen will indicate 0 Gbps for the Negotiated Link Speed. | Use the Filesystems or TCP/IP statistics pages to view the actual speed in Mbps.<br>This issue will be fixed in a future release. |
| | While servicing a CIFS file open request, the Slammer control unit having the VIF warm starts due to an oplock break failure. | Contact Technical Support.<br>This issue will be fixed in a future release. |
| | In rare circumstances, a request to modify a filesystem can fail if a system warmstart occurs while the request is executing. | Retry the failed request.<br>This issue will be fixed in a future release. |
| | When a large number of files are deleted while the system is busy, the filesystem can be tied up by the delete request, which may lead to the healthcheck detector to falsely assume the delete thread is hung and to warmstart the CU. | This issue will be fixed in a future release. |
| | On dual NAS systems, filesystem assignment to a given CU or Slammer is not possible. The system does filesystem assignments automatically. | This issue will be fixed in a future release. |
| | While using the supporttool wbinfo to enumerate local groups, the Pillar Axiom system might warmstart under certain circumstances. | Do not use the "wbinfo -g" option.<br>This issue will be fixed in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | Tape device names on a Pillar Axiom system may change during the course of a software update or a restart of the system, causing all tape and robotic operations to stop working. | Rename the tape devices in the backup application to reflect the changed names.<br>This issue will be fixed in a future release. |
| | When there is high volume CIFS traffic and multiple accesses to the same file with 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. |
| | The GUI displays the amount of space that belongs to a single snapshot (the amount of space that will be freed if a single snapshot is deleted). The display can show "0 GB", even though there is data associated with that snapshot. The GUI does not show the amount of space that belongs to more than one snapshot. That space is not freed if a single snapshot is deleted since it belongs to more than one snapshot. | If all the snapshots that contain the space are deleted, the space will show in the last snapshot that contains the space. The space and will be freed when the last snapshot that contains the space is deleted.<br>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.<br>This issue will be fixed in a future release. |
| | A deadlock can occur due to a race condition between moving files from subdirectories to parent directories while reading attributes of the parent directory, as in the following two actions:<br>`mv dir1/dir2/file1 dir1/file2`<br>`ls -l dir`<br>Such a deadlock causes a warmstart. | This issue will be resolved in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | On MAC OS X, if an NFS client using UDP protocol sends a request to a remote subnet and the Pillar Axiom system has a virtual interface on the same subnet as the client, an interaction with sendback routing may result in failure to mount. The UDP responses will not use sendback routing that returns the response on the interface on which it was received. | Use TCP for NFS mounts.<br>This issue will be resolved in a future release. |
| | Restoring a block backup to a target filesystem that is larger than the source causes the filesystem size to be set to the size of the source. | This space can be reclaimed by manually resizing the restored filesystem to the desired size.<br>This issue will be resolved in a future release. |
| | If an NFS exports file is uploaded and there are exports that modify existing entries, the modifications will not be effective. | When uploading an exports file, be sure there are no duplicate entries in the file and that none of the entries attempt to modify existing exports.<br>Remove any existing exports that require modification and enter the new information.<br>This issue will be fixed in a future release. |
| | When you copy a file from a CIFS share to a local drive, change its attributes, and copy it back to the share, the file properties page does not reflect the changes made in the local disk drive. | Move the file to the local drive instead of copying it. Change its attributes and then move it back to the share.<br>Another alternative is to reset the attributes.<br>This issue will be resolved in a future release. |
| | When the system is shutdown cleanly and then restarted, if a Brick is off line during the restart, filesystems on that Brick will remain offline after the restart completes, even if the Brick comes back on line. | Contact Technical Support for guidance in bringing the filesystems back online.<br>This issue will be resolved in a future release. |
| SAN | An error will occur if a user attempts to map a LUN to a host in the GUI by using a LUN number that the host is already using for another LUN. | View the LUN mappings for the desired host and choose an available LUN number for use on the LUN mapping page in the GUI.<br>This issue will be resolved in a future release. |
| | When restarting a Pillar Axiom system that has more than one Slammer and that has LUNs that are automatically assigned to Slammers, the LUNs may be Offline after the restart completes. | Restart the system again.<br>This issue will be resolved in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | Device Mapper may have difficulty in booting up if it encounters a LUN with corrupt vtable entries. | Un-map the LUN and reboot the server.<br>Upgrading to the latest version of the Linux Device Mapper and multipath tools with axiompmd (APM daemon) should resolve most situations. |
| | The pdscli command GetAllAvailableLUNNumbers shows a LUN number as available for the affected host, whereas the GUI does not show the LUN number as available. | In pdscli, choose another LUN number from the available LUN numbers.<br>This issue will be resolved in a future release. |
| | When a software update is performed under the following conditions, write failures may occur due to I/O timeouts:<br>• After a redundant LUN loses consistency.<br>• The mirror rebuild is in process.<br>• Writes are occurring on the LUN.<br>Depending on how the SAN host handles the write failure, data corruption could occur. | Queisce I/O before performing a software update.<br>Note: This issue does not apply to LUNs with standard redundancy.<br>This issue will be resolved in a future release. |
| | Performance degradation of non-optimized path accesses causes extent lock thrashing, resulting in a PANIC and a system warmstart. | This issue will be resolved in a future release. |
| | On SAN systems that are upgraded from one to two Slammers, the system may rebalance the LUNs to the new Slammer without prompting the administrator whether the rebalancing should occur. | To prevent automatic rebalancing, specify the Slammer assignment prior to adding the new Slammer.<br>This issue will be resolved in a future release. |
| | When a Clone LUN other than the oldest for a snapshot tree is deleted, the snapshot repository space that was allocated to the clone is reassigned to the next older clone. While this reassignment occurs, all older clones and the next newer clone (or the source if the newest clone is being deleted) are not accessible. This can take a long time when much data exists that has to be reassigned and may cause problems for applications and command timeouts within the system. | Whenever possible, delete only the oldest Clone LUNs rather than intermediary ones.<br>This issue will be resolved in a future release. |

| Slammer type | Pillar Axiom server issue or impact | Workaround or planned fix |
|---|---|---|
| | 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.<br>This issue will be resolved in a future release. |
| | Installing AxiomONE Path Manager 2.1 for AIX on a host system can reduce access performance from that host to Pillar Axiom LUNs. This performance reduction is a side effect of an error handling policy used by that release of APM. | Uninstall APM and use just a single path to access the LUNs.<br>Alternatively, contact Pillar Technical Support for help in disabling the error handling feature (which may result in some error situations not being handled well).<br>This issue will be resolved in a future release. |
| | SAN host continues to see LUNs from the Pillar Axiom after they are removed using the Pillar Axiom GUI and reconfigured back into the fabric. | Map the LUN to another host.<br>This issue will be resolved in a future release. |
| | In the Command Line Interface, when using PerformBackgroundLUNCopy to perform a background copy of a LUN, exceeding 12 simultaneous background copy operations may cause the Slammer CU to stop responding. | Do not initiate more than 12 simultaneous background copy operations. The limit of 12 simultaneous background copy operations will be enforced in a future release. |
| iSCSI | It is not possible to create multiple iSCSI sessions to Pillar Axiom systems through a QLogic 4052C iSCSI HBA when using the latest QLogic 4052C HBA software and the latest Microsoft iSCSI Initiator software. | Pillar Data Systems is investigating to determine the appropriate combination of QLogic and Microsoft software versions to use as a workaround for this problem.<br>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 are no longer product issues.

**Table 6 Resolved Pillar Axiom 500 server issues**

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| 03.00.00 | ALL | In the event that a Brick is under an extreme I/O load during a non-disruptive software update, it is possible for the RAID controller that is being upgraded to warmstart during the update. Although the administrator will see a RAID controller fault event, the controller will eventually restart with the new Firmware. During this event, users should not lose access to data. |
| | | In rare cases, the private out-of-band communication between the FC RAID enclosure and the FC Expansion enclosure can become inoperable. As a result, LED beaconing, environmental services, and the ability to perform software updates become inoperable in the Expansion enclosure. Data will, however, still be accessible. |
| | | The Pillar Axiom system fails to start up after multiple system panics caused by internal timeouts. |
| | | The 5-min timeout used during Call-Home transfers of log bundles by means of a proxy server is often exceeded, which causes the transfer to fail.<br>The fix increases the timeout to 2 hours. |
| | | In rare circumstances, a Slammer control unit (CU) will fail a health check (because of an internal problem), which results in the CU warmstarting. |
| | | The Pilot may become unresponsive, and an attempt to power cycle the Pilot results in continuous Pilot failure and rebooting. This may be more likely to happen after a software update or when a Pilot is having hardware problems. |
| | | In rare instances, the software upgrade process is not able to move uploaded software to the backup Pilot control unit. This failure results in the wrong versions of software displaying in the GUI. |
| | | When attempting to start the AxiomONE CLI, if the program is not present, the error message is not clear that the file is missing. |
| | | Sometimes when the system flushes I/O, it may not move the data to storage properly, causing the Slammer control unit to fail over and go Offline. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | If a Brick remains active on the Fibre Channel loop, but gets stuck in a state where it returns "busy" to certain operations, status updating on that Brick may be held up until those operations stop reporting "busy". As a result, the UI does not report a problem for the Brick even though the Brick in fact has a problem. |
| | | Insufficient information provided as to why the Verify step in Guided Maintenance failed.<br>The fix provides a “Load Bios Failed” Administrator Action when the Slammer PROM cannot be updated due to failures or compatibility issues. |
| | | Under Guided Maintenance, the Prepare step does not fail over the Slammer control unit (CU) being targeted for replacement. |
| | | Occasionally, when performing multiple backups to disk, timeouts may occur in the software, causing the system to lock up. |
| | | AxiomONE Capacity Planner does not give the user a choice of Premium for the performance priority, even when Fibre Channel bricks are present in the storage array. |
| | | Sometimes an attempt to modify the size of a logical volume (filesystem or LUN) will fail. |
| | NAS | Under certain conditions, a CIFS NTLM authentication failure may cause a warmstart. The failure occurs when the Pillar Axiom CIFS server is able to initiate communication with the domain controller, but the subsequent authentication steps fail (for example, when the machine account password is rejected). A configuration where the CIFS server has an active VIF on each control unit is more likely to encounter this failure. |
| | | In some circumstances, the filesystem over time leaks inode entries. When the inodes run out, the system normally recovers by a warmstart to free all inodes allowing system operation to continue. If, however, the inodes run out while the filesystem is being mounted and quota files are being accessed, a filesystem is incorrectly marked corrupted. |
| | | NAS Slammer control unit warmstarts when oplock break takes more than 5 min, which can occur, for example, because of heavy Brick loads. |
| | | Under certain circumstances, such as when a large number of users are doing many file operations, the Pillar Axiom may detect unfreed memory resources during the time of a filesystem unmount. In this circumstance, the system performs a warmstart to recover. |
| | | A NAS Slammer control unit (CU) warmstarts when a persistence operation initiated from a CIFS component is very slow (resulting from, for example, Brick Busy conditions). Warmstart is initiated only when the CU uptime is more than 1 hr. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | The NAS Slammer control unit with the VIF warmstarts when an NT_RENAME CIFS request with an information level of RENAME_FLAG_HARD_LINK(0x103) or RENAME_FLAG_COPY(0x105) is received from a CIFS client. |
| | | When a user maps a share on a Windows client that uses NTLM authentication, sometimes the Pillar Axiom system is unable to connect to the Domain Controller (DC) to authenticate the user, in which case the user receives the following error: "NT_STATUS_INVALID_COMPUTER_NAME".<br>This issue can occur after reconfiguring or upgrading the DC and when the Pillar Axiom system accesses the DC from multiple Slammer control units. |
| | | When trying to download a quota report, you sometimes may get a system error. |
| | | Oplock breaks may fail repeatedly when a client tries to open a file, leading to a Slammer control unit (CU) failure. After the CU recovers, the Oplock break will proceed. |
| | | Sometimes when the CIFS server accesses a Domain Controller from the Preferred Servers list to join a Domain, the request will fail with a “Join domain failure” event. |
| | | TCP/IP and VIF statistics are reported only for the control unit (CU) containing the primary VIF. The system leaves out these statistics for CUs containing secondary VIFs. |
| | | In very rare conditions, the filesystem metadata update missed a single block write causing the filesystem to declare corruption. |
| | | CIFS users may experience delays in mounting shares (mapping drives) when internal resources are not immediately available. Users have observed delays of 10 sec or more when Brick activity is very high (especially on NAS/SAN combination systems), or when many users attempt to connect simultaneously. The fix caches the required data and improves the response time. |
| | | If CIFS encounters a delay during Join Domain processing or Join Domain status queries, the system may warmstart. Delays could be due to DNS servers, NIS servers, or Domain Controllers that are not responding. |
| | | Scanning two or more WORM filesystems in parallel can consume all of the CPU resources in a Pillar Axiom system. This condition can cause the Slammer control unit (CU) to fail its health check, resulting in a failover and warmstart of that CU.<br>The fix schedules the work and allocates Slammer resources in a way so that the Slammer can respond to other requests. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | Under some NFS high workload conditions, I/O will stall and the NAS server will warmstart. This may occur when three conditions are true:<br>• The control unit servicing NFS (with the VIF assigned) is different from the control unit which owns the filesystem (thus filesystem traffic is active between the nodes).<br>• The NFS server workload is sufficient to be consuming all internal NFS resources.<br>• The NFS clients are intermittently closing the TCP connections. |
| | | The filesystem runs out of internal memory during lock manager "unlock all" operations if the number of locks are substantial. |
| | | In rare circumstances, when a filesystem is taken offline, data access to the control unit may seem to lock up. Data access appears hung until the unmount finishes. |
| | | CIFS may not shutdown promptly under certain conditions, causing a healthcheck and a warmstart. If a shutdown request is received while CIFS is attempting to communicate with DNS or NIS servers that are not responsive, the shutdown may be delayed and fail. |
| | | In rare situations, the system miscalculates the outstanding I/O count for a filesystem, causing an indefinite I/O wait condition eventually leading to a warmstart of the control unit. |
| | | The system behaves poorly when a CIFS request uses more than 128 network packets, often resulting in the Slammer control unit restarting. This restart breaks all CIFS connections.<br>The fix enables the system to correctly handle CIFS requests even when the system receives a request that uses hundreds of network packets. |
| | | Creating a SecureWORMfs filesystem by means of the Configuration Wizard will generate an error. |
| | SAN | SAN runs LUN configuration commands through the task set. In this case, the configuration command could not run due to an ACA in the task set. |
| | | For a non-owner command that returned a check condition before being enabled, the system does not wait for the message received from the owning control unit before completing. |
| | | In rare circumstances, the system will Segfault with two SAN Panics, which cause the system to restart. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | If a LUN is configured to have access bias set to sequential and I/O bias set to write, the system sets the caching strategy to write-through. If either of these attributes is then set to a different value, the LUN's caching strategy would be incorrectly set until a cold start was performed.<br>The code for setting of caching mode has been extensively reworked to preclude this issue. |
| | | Migration of LUNs from SATA Bricks (Archive through High Priority) to more than two FC Bricks (Premium Priority) may result in a less-than-optimal distribution over the target Bricks.<br>Also, a "Temporary" Priority might be set that makes the LUNs appear to have the same performance level after migration. |
| | | After updating to Release 2.6 (or a later 2.x version), periodic PilotInternalError events are generated and accompanied by Call-Home events. |
| | | While running FSCK, the active Pilot control unit (CU) encounters a deadlock that would normally cause the Pilot CU to fail over. However, another deadlock prevents the CU from completing the fail over, so instead it just hangs. |
| | | A Slammer control unit could not fail back because the Slammer software attempted to configure a LUN that had previously been deleted. |
| | | The SAN host can see LUN mappings that are not visible to the Pillar Axiom software. This lack of synchronization results in not only unmapped LUNs being seen by the host but those LUN numbers being unavailable to future mappings to that host. |
| | | The Modify LUN Mappings page of the GUI shows the following for broken LUN mappings:<br>HostName = blank, LUN Number = empty, Port Mask Set = Unset.<br>The mapping, however, is available and can be seen at the host. |
| | | If during a Copy LUN process the background task invoked by the Copy is cancelled, the resulting LUN displayed in the GUI is a Clone LUN, which was used to begin the Copy LUN. If data is written to this LUN, it will be deleted once the repository space is full. |
| | | An error may occur when attempting to map a LUN to a SAN host using a LUN number that has already been used by that host for another LUN. |
| | | If you attempt to delete a LUN that has associated Copy LUNs [Clone LUNs], the error message returned incorrectly indicates that the expected failure to delete the LUN was due to an incomplete Volume Copy.<br>The fix returns the correct error message to distinguish among the different cases. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | If Clone LUNs have been created for a LUN and the parent LUN is modified to disable the SnapLunSpace, the system will return the error: “Some of the values you entered are invalid. Please correct them and try again.” |
| 02.08.00 | ALL | File data may get compromised when the I/O load on a Brick storage enclosure is extremely high. |
| | | After making changes to the priority settings of a LUN or filesystem, subsequent changes to the same LUN or filesystem result in an "Internal error" message, and the latest change does not take effect. |
| | | When replacing a fan module under Guided Maintenance, if the old and new fan modules are swapped too quickly (less than 5 sec), the system may not have enough time to sense that the new module has been inserted and run diagnostics properly. In such a case, the system reports that the new fan has failed.<br>The fix allows the fan module to be detected. |
| | | After saving a LUN configuration created in the AxiomONE Capacity Planner, the saved configuration cannot be used to create the actual configuration. |
| | NAS | When a Windows domain controller locks out a CIFS user account (for an unknown reason), the client retries to authenticate the locked out user several times. Sometimes this activity causes the client connection to reach its imposed memory limit. When the limit is reached, the affected Slammer control unit (CU) warmstarts.<br>The fix terminates the client connection to avoid a Slammer CU warmstart. |
| | | When the preferred server is set in the CIFS Server Comment field by specifying an IP address, the "set kerberos password" request does not contact the KDC host specified but instead depends on the server address returned by the DNS server.<br>The fix is to modify the set kerberos password code path, so that it overrides the KDC host returned by the DNS server with that read. |
| | | When copying a filesystem and one of the internal tasks hangs (a background task in the GUI or CLI task response is not proceeding), if you attempt to stop the task, the active Pilot control unit will restart.<br>Note: Be patient while copying large filesystems, as such an operation takes a while. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | A File Server can hold 256 unique client host names for the purpose of lock reclaim notification. Some clients (such as PC-NFS and DHCP) that occupy the lock reclaim table may never return to clear the locks during the grace period after a system restart. If client locks are not reclaimed during the grace period, the Pillar Axiom system may not clear those un-reclaimed locks properly when the grace period after a system restart expires, which can cause the host limit to be exceeded.<br>The fix for Release 02.08 properly clears these locks when the system is restarted. (A future release allows the stale locks to be cleared by an operator action.) |
| | | When a Pillar Axiom system has joined a Windows domain consisting of multiple Domain Controllers and those controllers become unreachable, adding a new VIF operation which is part of the server configuration process will cause the owning Slammer control unit to warmstart. |
| | | On Pillar Axiom NAS systems running under extremely heavy loads, the TCP/IP stack software may (by design) run out of buffer resources to continue processing I/O, resulting in a software deadlock, which eventually causes the Slammer control unit to warmstart. |
| | | The CIFS domain user on the client was locked out on the Domain Controller. When this occurs, the CIFS client repeatedly tries to authenticate the CIFS user at a rapid rate. All these authentication requests are rejected. This situation consumes a large amount of memory and causes the connection to exceed its self-imposed memory budget, resulting in a warmstart.<br>The fix simply closes the connection. |
| | | The share name "homes" may map to a different filesystem than the one configured in the share definition. |
| | | Sometimes internal house-keeping tasks on the Pillar Axiom system consumes all available CPU and prevents any user-initiated operations from completing. In such a case, the system may fail to process protocol requests for up to five min, after which the system warmstarts. |
| | | During CIFS I/O, when an SMB signature verification fails, attempts to end the CIFS connection will fail and cause a software recovery. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | SAN | Under some scenarios during Auto Contingent Allegiance (ACA) compliant error handling, valid transfer to a LUN may be blocked, causing an internal SAN software failure.<br>This situation occurs when an error is returned on I/O transfer and, while processing this error, an ACA occurs, which can result in the blocking of on-going commands to the same LUN. Because we had overloaded a variable to handle I/O errors and blocking, blocking the command caused us to mishandle the data handshake, resulting in the panic.<br>The fix treats these conditions separately. |
| | | In very rare circumstances, if persistent loop instability is detected between the Bricks and the Slammers, a race condition developed where ports could end up doubly counted, causing the system to restart. |
| | | If the Administrator, using the GUI, modifies a LUN to change either the LUN's Slammer CU assignment or the LUN's Relative Priority or Redundancy, without having first viewed the (LUN) Mapping tab, the mappings for the modified LUN will be removed. |
| | | The SAN host issued a SCSI Mode Select command to change some configuration settings of a LUN. These settings are written to both CUs of a Slammer. When syncing these changes, the other CU warmstarted, preventing a return response to the sync, which after two min caused a health check to be generated. LUN configuration is automatically resynced on restart, so waiting for the return was unneccessary.<br>The fix removed the wait for completion. |
| | | In cases where there is an extremely unstable Fibre Channel link (in other words, multiple link downs in succession), it is possible to not clean up a task that has been aborted by the host during link down processing.<br>The task is missed because we have not identified the sending host at the time of the link down. This task will begin to run if the host that sent that I/O logs back in after link up. When we try to transfer data, the Fibre Channel chip no longer recognizes this I/O, and we fail.<br>The fix aborts all tasks properly. |
| 02.07.00 | ALL | Pillar Axiom systems updated to release 02.06.00 may receive messages indicating that the packaged RAID firmware is not compatible with the firmware installed on the system. In this case, the RAID firmware on the system is not updated. |
| | | For Kerberos/ADS authentication, map drive requests could fail occasionally, or the system could warmstart. This is due to a memory leak in the Kerberos/ADS authentication code path. The fix eliminates the memory leak. |
| | | A bad disk drive was unable to be properly initialized after an insertion. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | In rare circumstances, an I2C bus on the FC Extension Brick could get stuck, limiting communication to the ES module in that enclosure. |
| | | In very rare cases, the seed value used to uniquely identify communications between the FC RAID and FC Extension controllers could duplicate, causing ancillary communications to fail. |
| | | When a Brick is removed from the system, the Pilot's internal database may be incorrectly updated. If another Brick is added to the system later with the same name as the Brick that was removed, the new Brick will not be added correctly (RAID burn will fail). Also, CLI commands such as GetAllBricks and GetBrickDetails will show "FQN not found" where the new Brick FQN should be shown. |
| | | After a software fault, in rare circumstances, a Slammer control unit could hang and not restart properly, requiring the CU to be power cycled to recover. |
| | | There was a very small timing window during a single controller restart where the data being sent between the two controllers could get sent before the restarting controller was ready. |
| | | Under unusual circumstances when a control unit (CU) is servicing IOs for the same VLUN/range that its buddy CU is servicing, the IO might become starved, which causes a warmstart. |
| | | There is a very small window in the Brick removal process where the system might not notice the removal. |
| | | Bad hardware on a RAID controller caused the partner controller to be reset during replacement. |
| | | Under certain circumstances, incorrect access control of a shared internal data structure could lead to its content being compromised, resulting in a software panic and recovery. |
| | | On a system that experiences CU failovers, a request to Persistence that is in process might not complete, causing the request to hang and fail an operation. |
| | | In rare circumstances, an I2C bus on the JBOD enclosure could get stuck, limiting communication to the ES module in this enclosure. |
| | | When manually collecting logs, the full set of live Slammer logs was not present. |
| | | In rare circumstances, FC loop instability can cause a command to hang in a RAID controller. The Brick error handling properly detected this occurrence and warmstarted the Brick to recover the system. |
| | | Sometimes, when a LUN or filesystem goes offline, the system may segfault on pending quota operations. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | Under certain situations, the task bar shows a Call-Home task being executed that never completes. In these cases, a Call-Home transfer is stuck, and from this point forward, will wait five min before aborting a failed transfer and failing the Call-Home. |
| | | If a Slammer fan module is replaced using Guided Maintenance, the system erroneously puts the LUNs and filesystems into conservative mode. |
| | | An erroneous PCO record in the Pilot's database causes periodic PilotInternalError warning events to be generated, which show up in the Event Log. CallHome actions and administrator emails are triggered as well. |
| | | On a restart due to a system core dump, the most recently generated core dump on the passive Pilot will not be detected. |
| | | Under rare conditions, during a Brick software update, a disk drive may be incorrectly taken offline due to a hardware disconnect. |
| | | The Pilot has a service installed called postfix. Although not in use, this service can cause the Pilot disk drive to fill up with log files, resulting in a Pilot restart or update problem. The fix removed this service. |
| | | If too many locks (> 251) are released at once (for example when a client holding many locks disconnects), the system warmstarts. |
| | | After a software fault, a Slammer control unit may hang and not restart properly, requiring the CU to be power cycled to recover it. |
| | | A data LUN reappeared after background copies had been cancelled. This could happen if, for example, a Brick shows up.<br>The fix now starts background copies asynchronously, so that separate tasks are started, and the existence of background copies is no longer masked by the HandleAppearingLUNTask. |
| | | In rare cases a code bug could cause command execution times to get very large. |
| | | Sometimes FC Bricks can be lost after a power cycle that was not preceded by a controlled shutdown. |
| | NAS | NFS Linux clients may get stuck in repeatedly calling readdir/readdirplus. This can happen when the initial call is followed by adding files to the directory, and then another readdir/readdirplus call. The fix corrects the system logic to handle readdir cookies in this code path. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | Map drive requests fail intermittently, either timing out or with a Kerberos time skew error. Longer delays may cause the system to healthcheck and warmstart. This behavior is more likely to happen on a busy system that is processing many authentication requests. The fix provides more equitable response times to queued requests. |
| | | Slammer control unit might warmstart under a heavy load during a DNS lookup request. |
| | | Mapped drive fails intermittently on Windows clients, accompanied by the following error message:<br>“The mapped network drive could not be created because the following error has occurred:<br>“There are currently no logon servers available to service the logon request" |
| | | NFS segfaults and warmstarts when it receives a zero length RPC request. The fix makes the code more resilient to unexpected client requests. |
| | | CIFS warm starts occasionally when it cannot allocate an internal shared resource. |
| | | When a File Server is configured to use Active Directory with Kerberos authentication and a CIFS client maps a drive to a filesystem, the system may warmstart when an application retries Kerberos authentication many times due to some failure condition. |
| | | The CIFS server may segfault and warmstart when a client sends a TreeDisconnect request after a LogoffAndX. The CIFS server discards the session credentials during the LogoffAndX, thus they were unavailable to handle all potential disconnect/close processing (in this particular case, a file access time update). The fix allows this sequence of requests. |
| | | When running dbench (an open source program to test filesystem access), sometimes the Axiom CIFS server warmstarts. This occurs when the connection is being closed and the system cannot break an oplock. |
| | | CIFS warmstarts occasionally when the customer's primary DNS server is offline.<br>The fix allows CIFS to be more resilient and fail individual requests without warmstarting the system. |
| | | CIFS segfaults and warmstarts intermittently during a directory delete operation. This may occur when another application is creating files into the directory while a delete is in progress. The fix eliminates a double free operation on a specific error codepath. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | Map drive requests were rejected from a Citrix server. Citrix uses a single connection to share map drive requests (TreeConnectAndX CIFS commands), and an internal limit per connection was reached. The fix removes the artificial limit. |
| | | VPN user authentication fails due to a consistent checksum error from the Pillar Axiom system (CIFS DCERPC failure). The CIFS domain name is used within the RPC request, but when the first component of the fully qualified domain name does not match the NetBIOS domain name, the request is rejected for a checksum error. The fix retrieves the correct NetBIOS domain name from the customer's LDAP server. |
| | | When the VIFs of a CIFS server span both control units of a Slammer, the domain controller may reject concurrent NTLM authentication requests and cause map drive requests to fail intermittently, which could lock out a user's account. This condition could also lead to a memory leak condition on the NTLM authentication error path.<br>The fix synchronizes NTLM authentication requests when the CIFS server spans both control units. |
| | | CIFS segfaulted and warmstarted when a client sent an oplock break request after a LogoffAndX. Because the CIFS server discarded the session credentials during the LogoffAndX, the credentials were unavailable to handle the oplock break processing. The fix allows this sequence of requests. |
| | | A Slammer control unit might warmstart under heavy loads during a DNS lookup request. |
| | | When PerformFSCKTask and CreateScheduledFileSystemSnapshotTask run at the same time, the active Pilot control unit may fail over to the passive Pilot CU. |
| | | The count of blocks associated with a file goes out of sync when a file is truncated. |
| | | CIFS segfaults and warmstarts when an unsupported RPC request is received. A client may have used MMC to query the number of active CIFS sessions and sent the RPC function SRVSVC:NetSessEnum. This RPC request is not currently supported in the Pillar Axiom CIFS product. The fix returns an appropriate error code. |
| | | CIFS segfaults and warmstarts when an RPC pipe authentication request occurs before a SessionSetupAndX session authentication request. The fix allows this sequence of requests. |
| | | When multiple NTLM authentication requests fail due to a problem contacting the domain controller, CIFS can exceed its internal resource limits for the File Server and result in a Slammer warmstart. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | Under some conditions, CIFS may exceed its internal resource limits for the connection while waiting for an oplock break. This causes a segfault and warmstart. |
| | | CIFS exceeded its internal resource limits for the connection. This caused a segfault and warmstart.<br>The fix has CIFS return an appropriate error code to the client. Future releases will increase the internal resource limits. |
| | | When a filesystem cannot be fixed, a “Filesystem Check (FSCK) Failed” Administrator Action is generated. If the user doesn’t open the Administrator Action, the user won’t know there is a choice to restore the filesystem from a snapshot.<br>The fix changes the title of this admin action to “Filesystem Check (FSCK) Completed” and changes the text within the Administrator Action to make it more consistent with other actions. Also, a warning was included with the Delete choice. |
| | | The system warmstarts when it exceeds the memory allocation limits for internal CIFS structures. In certain conditions, CIFS Tree Disconnect requests are rejected and active Tree Connections are not freed. This can occur when a Tree Disconnect follows a LogoffAndX, and the user ID changes prior to the Tree Disconnect request. |
| | SAN | During SAN host discovery by Pillar Axiom systems equipped with a 4 GB SAN controller, the Slammer could panic with the message "PANIC_ASSERT: hash lookup failed" and then warmstart.<br>The fix implemented a new hashing algorithm to alleviate this issue. |
| | | During SAN host discovery by Pillar Axiom systems equipped with a 4 GB SAN controller, the Slammer could panic with the message "Hash lookup failed for <portid> attempts 3". This would sometimes lead to a Slammer warmstart.<br>The fix implemented a new hashing algorithm to alleviate this issue. |
| | | On the Storage>Hosts page in the GUI, selecting the “Associate a Host’” action causes the active Pilot control unit to fail over to its partner CU. The page will not be displayed. |
| | | For Pillar Axiom systems having one or more SAN hosts running APM, occasionally a Pilot control unit will unexpectedly fail over to the partner CU. |
| | | If a LUN is configured for Sequential access with a Write IO bias, the system sets the caching strategy to write-through. If either of these parameters is changed later, the LUN's caching strategy is not set to write-back, which it should be. |

| Fixed in Release | Slammer type | Pillar Axiom server issue or impact |
|---|---|---|
| | | System fails when deleting non-eldest Clone LUN.<br>Refer to Sections 9.20 and 9.21 for related information.<br>Release 02.07.00 contains an incremental improvement only and allows for slightly better handling in this scenario. A full resolution involving changes in design and implementation will come in a future release. |
| | | When you click on Storage>LUNs in the GUI, the GUI response is slow, especially when the Pillar Axiom system has a high-volume configuration of LUNs. |
| | | When a data migration setup is interrupted and the Pillar Axiom system is then updated to new code that doesn't expect migrations to be in that state, the system will not advance the setup of the migration to move beyond the partial setup. The DeleteLUNTask safely refused to try to delete a LUN that was in this state of partial setup.<br>The fix recognizes the old state and advances the data migration setup in the cold-start task so the data migration can continue. |
| | | If the system sees a pending delete SAN LUN or Clone LUN when restarting, the system could mistakenly attempt to configure the LUN. Doing so would cause the Slammer to repeatedly panic. |
| | | If JavaScript is not enabled when using the GUI, operations such as CreateLUN and ModifyLUN can erroneously allow you to specify a LUN number even for LUNs that are mapped. The GUI will display the error message "A system error has occurred: Internal UI error occurred". The LUN creation or modification will not happen and you will be taken back to the previous page. |
| | | For LUNs greater than 810 GB, you cannot create shadow copies of those LUNs if the Clone LUN capacity has not been previously allocated on the LUN. |
| | iSCSI | When an iSCSI host is shutdown, the APM field correctly displays "Not communicating" whereas the iSCSI host port status incorrectly displays "Connected" on the hosts page of the Pillar Axiom GUI.<br>The system now properly shows the connection status of an iSCSI host when it is powered off. |
| | | In an iSCSI session, when Login C attempts to reinstate Login B, which is attempting to reinstate Login A, Login C panics.<br>The fix serializes iSCSI session reinstatements when the system receives duplicate logins in parallel. |
| | | The host entry for the iSCSI IQN would go away and no longer display on any of the host display pages. In addition, the Storage>Hosts page would not display the iSCSI IQN data for the new owning APM host. However, the Initiator is associated with the APM host and is accessible through it. |

# 9 Additional Notes

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

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

## 9.1 Host Queue Depth on SAN Hosts

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

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

## 9.2 Limitation of the Linux sginfo Utility

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

## 9.3 LUN Ranges in Windows

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

## 9.4 Issues with LUN Capacity Calculations on Solaris

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

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

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

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

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

## 9.5 VMware ESX Server 3.0.1 Connections to Pillar Axiom iSCSI Systems

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

## 9.6 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.

## 9.7 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.

## 9.8 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.

## 9.9 Blacklisting Local Disk Drives on RHEL4 Platforms

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

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

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

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

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

## 9.10 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 Technical Support 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.

## 9.11 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.

## 9.12 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.

## 9.13 Non-Disruptive Software Updates

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

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

**Important!** For Pillar Axiom systems that have one or more NAS Slammers, try to quiesce all I/O to the system before performing a non-disruptive software update.

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

However, any application or operating system that bypasses the Fibre Channel protocol stack and issues SCSI commands with short timeouts may not be capable of handling the brief interruption of a non-disruptive software update. For example, the Microsoft Cluster Service will initiate retries in 3 seconds and, at 7 seconds, will begin reconfiguration and recovery that will probably disrupt any applications using the Cluster Service. Microsoft, however, is promising to fix this issue in the upcoming release of the Windows “Longhorn” server (successor to Windows Server 2003).

## 9.14 WWN Designation Changed in Release 2.0

Starting with the 2.0 release, 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.

## 9.15 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.

## 9.16 Status of Filesystems and LUNs

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

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

## 9.17 Change in Default Failback Configuration of NAS Slammers

Automatic failback of NAS Slammers is the default configuration beginning with Release 02.00.00. 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.

## 9.18 Avoid Setting the Time by Means of Pillar Axiom Facilities

If the time on a Pillar Axiom system is being controlled by a time server, the date and time should *only* be changed on that time server.

If you are not using an NTP server, we recommend not changing the date by means of the GUI or CLI 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 Technical Support for recovery assistance.

## 9.19 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.

## 9.20 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.

## 9.21 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.

## 9.22 Reassigning a SAN LUN to another Slammer or Control Unit

If you reconfigure the Slammer or the Slammer control unit (CU) to which a SAN LUN is assigned, the Pillar Axiom restarts the LUN at the new location. As a result, users may experience a momentary loss of access to that LUN, possibly disrupting some applications. SAN hosts should be able to re-establish LUN access automatically. However, Pillar Data Systems recommends that you change LUN ownership only when I/O is not occurring.

## 9.23 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.

## 9.24 Running Slammer Diagnostic Tests from the GUI

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

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

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

If the diagnostic passes:

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

**Important!** Do not attempt to run Diagnostics on more than one Slammer CU at a time. Doing so may cause additional resources to go offline.

## 9.25 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.

## 9.26 GUI Accurately Displays the Status of Components

In earlier releases, during system or core restart, all components would show a status of “Booting”. Beginning with Release 02.00.00, the GUI more accurately displays the status of the components as they are discovered and initialized by the management software. For example:

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

## 9.27 A Disk Drive May Display Blank Data in the GUI

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

## 9.28 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 Technical Support for assistance in recovering.

## 9.29 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 Technical Support. Also, when you need to replace a Pilot CU, contact Technical Support for assistance.

## 9.30 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.

## 9.31 Replacing a Disk Drive

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

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

  **Important!** You should contact Pillar Technical Support before pulling a disk drive.

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

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

Contact Technical Support for a new replacement disk drive.

## 9.32 Moving Disk Drives

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

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

## 9.33 Reseat Disk Drives before Powering On New or Replacement Bricks

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

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

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

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

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

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

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

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

## 9.34 Testing Multiple Disk Drive Failures

**Important!** Do not test multiple disk drive failure scenarios in the same Brick storage enclosure without contacting Technical Support for guidance.

## 9.35 ACT LEDs on Disk Drives Can Blink When Inactive

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

## 9.36 Replacement of Brick Storage Enclosures

To avoid data loss, contact Technical Support before you attempt to replace an entire Brick storage enclosure or Slammer storage controller. Technical Support can help you determine whether a particular filesystem or LUN is physically on the Brick.

## 9.37 Adding a Brick Generates Error and Warning Messages

When you add a Brick storage enclosure to an existing Pillar Axiom system, the system begins the process to bring the Brick online. While the system is bringing the Brick online, you may see a series of error and warning messages similar to these:

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

These messages are normal and to be expected. During the bring-up process, the status of the Brick will go from red to yellow to green. After the system completes the process, the Brick will show a Normal status and will remove all Administrator Actions related to adding the Brick. If any Administrator Actions remain, contact Technical Support.

## 9.38 FC Bricks in SATA/FC Systems Support Only the Premium Band

When adding Fibre Channel (FC) Bricks to a SATA system, the FC Bricks can only be used for the Premium priority band. This means, for example, that if you have a mixed SATA/FC system, the maximum possible space for the Archive performance band will be the total amount of SATA disk space you have. It also means that the FC space on the system cannot be used for any performance band except the Premium band. In other words, *bleed up* of lower-performance bands from SATA to FC storage in a mixed SATA/FC system is not supported in this release.

## 9.39 Testing RAID Rebuild and Simulated Drive Fault

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

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

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

If a drive is genuinely faulted, use the Guided Maintenance menus to Identify the Brick, note the position of the drive in the Brick, Prepare the System for replacement, and then "Replace Hardware" to remove the old drive and replace it from spares as instructed.

## 9.40 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.

## 9.41 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.

## 9.42 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.

## 9.43 Uploading Empty Files Not Allowed

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

## 9.44 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.

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 Pillar Technical Support 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 Pillar Technical Support for assistance in determining candidates for good snapshots before reverting in this manner.

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

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

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

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

1. Takes the filesystem offline so that no further changes can be made to the data.
2. Generates an Administrator Action that provides multiple options to the administrator:
    - Perform a filesystem consistency check.
    - Delete the filesystem.

- Revert to a previous snapshot.
  If you choose this option, all changes to the filesystem between the time of that snapshot and the time the filesystem issue was detected is lost.

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

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

   Note: If no good Snap FSs exist and the initial FSCK fails to recover the filesystem, it is recommended that you *not* delete the filesystem but instead contact Technical Support. Technical Support may be able to help you recover the filesystem.

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

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

## 9.45 Journal Data for SecureWORMfs Filesystems Always in Physical Storage

For compliance filesystems using SecureWORMfs, the journals for those filesystems are always kept on physical storage and the battery-backed memory is intentionally disabled. This feature does cause a small loss in performance but is necessary for data integrity.

## 9.46 Replication of a SecureWORMfs Filesystem

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

## 9.47 Increasing Redundancy Requires More Space Than Requested

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

## 9.48 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.

## 9.49 Number of Directory Quotas

The current limit to the number of directory quotas that can be managed is 100,000 per filesystem.

## 9.50 CIFS Support

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

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

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

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

Additional CIFS features will be supported in future releases.

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

## 9.52 DHCP Behavior on the Pilot

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

- Dynamically assigns only the public IP address of the Pilot.
- Locks the two private IP addresses.
- Retains DHCP settings during a Pilot failover.
- If the IP address is updated through 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.

## 9.53 Changing Slammer Port IP Addresses from Static to Dynamic

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

## 9.54 SMI-S Should Not Be Enabled on 512 MB Pilots

**Important!** By default, all 512 MB (older, non-RoHS) Pilots have SMI-S disabled on startup. Because of limited memory, you should not enable SMI on these Pilots.

## 9.55 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.

## 9.56 Disabled/Excluded Slammer States

Repeated failure of a Slammer control unit (CU) can result in that CU becoming non-operational. If the repeated failures occur during normal operation of the Pillar Axiom system, the system marks the CU as Disabled; if the failures occur during startup, the system marks the CU as Excluded. This behavior may be triggered by repeatedly testing power failure simulation for Slammers.

If the system completes startup successfully, it will attempt to remove each Excluded Slammer control unit (CU):

- After startup completes, the CU should transition to Failed Over.
- If the CU is part of a SAN Slammer, the system should then attempt to transition the CU to Failback as long as the buddy CU on that same Slammer is online. If the Failback succeeds, the system puts the CU Online.
- If the CU is part of a NAS Slammer and automatic failback is *not* enabled, the system attempts to transition the CU to Failed Over and leave it there for recovery.
- If automatic failback of NAS Slammer resources *is* enabled, the system attempts to transition the CU from Excluded to Failed Over to Failback and, if that succeeds, the system puts the CU online.
- If the CU is not detected, but the other CU is active, a missing CU may show a status of Failed Over when it is really Offline, as in powered down or not connected to the private management network.

## 9.57 Hardware Lockout Due To Repeated Power Cycling

IMPORTANT! Do not repeatedly power cycle a Pillar Axiom system or any of its hardware components. Doing so may automatically trigger the hardware-fault lockout mechanism. The current thresholds for repeated power cycles are:

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

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

Contact Technical Support for assistance in recovering and restoring the components to service.

## 9.58 Powering Off a Pillar Axiom System

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

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

## 9.59 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.

## 9.60 Battery Insertion

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

## 9.61 Slammer Warmstart and Startup Failure Handling

Slammer 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 Pillar Technical Support for recovery assistance for any Slammer CU that is Excluded or Disabled.

## 9.62 Preventing Filesystem Corruption Due To Multiple Component Failures

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

## 9.63 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 Technical Support.

## 9.64 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.

## 9.65 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.

## 9.66 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.

## 9.67 NDMP DMAs May Disable Tape Storage Devices During a Software Update

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

## 9.68 Time Zone Must Be Reset after a Software Upgrade

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

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

## 9.69 Point-in-Time Volume Copy Space Allocation

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

## 9.70 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.

## 9.71 Link Aggregation over Fast and Gigabit Ethernet Switches

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

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

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

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

## 9.72 Vulnerability Scanners May Report False Positives

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

## 9.73 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 Technical Support.

## 9.74 Login to Oracle EM Plugin Fails

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

## 9.75 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.

### 9.75.1 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.

### 9.75.2 Behavior

#### 9.75.2.1 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.

#### 9.75.2.2 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.

#### 9.75.2.3 Using PCSH to Move an Initiator from an Associated Host to an APM Host

When an Initiator is moved from an Associated host to an APM host by a PCSH(V2) request, the mappings of the previous owner Associated host are moved to the new 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.

#### 9.75.2.4 Using PCSH to Move an Initiator from an APM Host to another APM Host

When a Initiator is moved from an APM host to another APM host by a PCSH(V2) request, no mappings belonging to the original owning APM host are moved to the new APM host. The old mappings to the Initiator are removed and the Initiator is mapped to the configuration of the new owning APM host.

#### 9.75.2.5 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.

#### 9.75.2.6 Notes

##### 9.75.2.6.1 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.

##### 9.75.2.6.2 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.

#### 9.75.2.6.3 Renaming an APM Host

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

# 10 Technical Documentation Errata

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

## 10.1 *Pillar Axiom Administrators Guide*

- The following features are discussed in this book but are not implemented in Release 03.00.00; they are however planned for an upcoming point release:
  - Oracle ASM performance profile (discussed on pages 33, 47, and 48)
  - Wide Stripe (1 MB strip) configuration option (discussed on pages 47 and 48)
- On page 52, under Step 3 of the Define LUN Quality of Service (QoS) Attributes task, please add the following notice:

  **CAUTION**: Under HPUX (and possibly other operating systems), changing the capacity of the LUN without first migrating data from the LUN will result in data loss. Storage Administrators need to be aware of how their operating system will behave when LUN capacity is increased, which may depend on the details of how they have configured their environment.
- Some 32-bit operating systems cannot access or manage NFS or combined NFS/CIFS filesystems of 2 terabytes or greater in size. If in doubt, check the operating system limits for your client and filesystem. For mixed 32-bit and 64-bit clients, limit filesystem sizes such that both clients will be able to work with those filesystems.
- The Guided Maintenance menu in the AxiomONE Storage Services Manager (GUI) incorrectly implies that you can replace Brick and Slammer chassis. Replacing Brick and Slammer chassis are not supported in this release.
- On the Resolve Connectivity Trouble page (page 257), the examples shown in the Syntax column in Table 16 should distinguish between syntax, environment variables, and examples. The commands that include examples in the Syntax column should be as follows:

| Command | Description | |
|---|---|---|
| `ping` | Command: | `ping -c 5 198.168.1.2` |
| | Environment Variables: | |
| `route` | Command: | `route show` |
| | Environment Variables: | |
| `arp` | Command: | `arp -a` |
| | Environment Variables: | |
| `ifconfig` | Command: | `ifconfig` |
| | Environment Variables: | `SOCK=`*`node_ID`* |

## 10.2 *Pillar Axiom Architecture Overview*

The following features are discussed in this book but are not implemented in Release 03.00.00; they are however planned for an upcoming point release:

- Oracle ASM performance profile (discussed on page 25)
- Wide Stripe (1 MB strip) configuration option (discussed on page 25)

## 10.3 *Pillar Axiom CLI Reference Guide*

Once the `PerformEnableSMIStartup` command has been executed, the Pilot needs to be restarted to enable the SMI Provider. Restart the Pilot using the following command:

```
$ pdscli submit -u <user> -p <password> -H < host-name-or-IP>
PerformRestart
```

When the system comes back online, the smProvider service should be enabled. Once enabled, the service will remain on through power cycles until you manually disable it with PerformDisableSMIStartup.

- **Important!** By default, all 512 MB-based Pilots have SMI-S disabled on startup. Because of limited memory, you should not enable SMI on these Pilots. For anything other than testing purposes, contact your Pillar Account Representative for a memory upgrade to your Pilot.
- The “perf” command is missing in Table 23 on page 542. To run this command, specify perf in the Command option of PerformSlammerCommand:

  ```
  $ pdscli submit -u <user> -p <password> -H < host-name-or-IP>
  Command=perf
  ```

  This command returns the CPU statistics for the idle process, the kernel time, and all process IDs that are running.

## 10.4 *Pillar Axiom NDMP Integration Guide for NAS Systems*

- On pages 15, 16, 17, and 21, change all <*pds-hostname*> references to <*pilot hostname or IP*>.