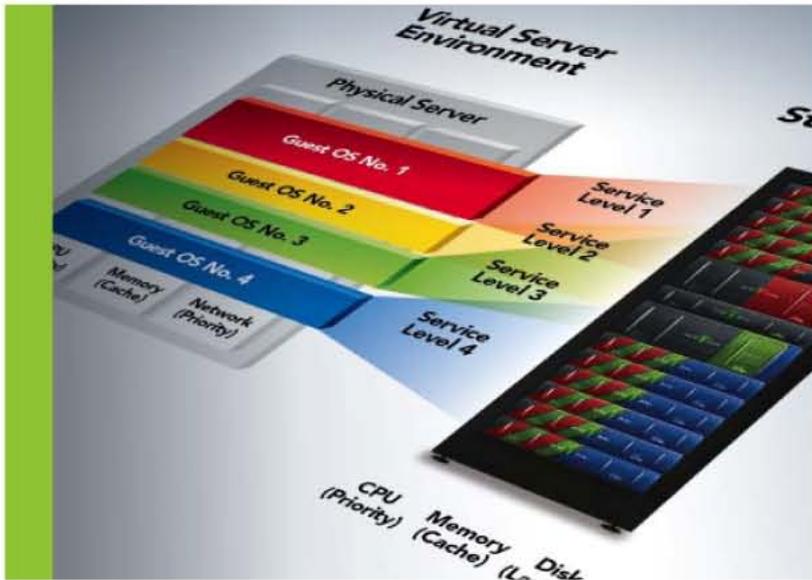


Pillar Axiom 500 and 600



Customer Release Notes

Release 5.4

ORACLE

PILLAR AXIOM

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

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

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

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

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

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

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

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

Copyright © 2013, Oracle et/ou ses affiliés. Tous droits réservés.

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

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

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

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

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

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

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



PILLAR AXIOM

Pillar Axiom 500 Pillar Axiom 600	Document Number: E-39219-02
	Document Title: Customer Release Notes

Revision History

Rev description	Effective date
Release 5.4 Rev 1.1	2013-02-15 [Pillar Axiom 500 and 600 SAN only]
Release 5.4	2013-01-31 [Pillar Axiom 500 and 600 SAN only]
Release 5.3	2012-03-15 [Pillar Axiom 500 and 600 SAN only]
Release 5.2	2011-10-01 [Pillar Axiom 500 and 600 SAN only]
Release 5.0 GA Rev 1.1	2011-06-21 [Pillar Axiom 500 and 600 SAN only]
Release 5.0 GA	2011-06-13 [Pillar Axiom 500 and 600 SAN only]
Release 4.3.1 GA Rev 1.1	2011-02-11
Release 4.3.1 GA	2011-01-19
Release 4.3.0 Beta	2010-12-28
Release 4.2	2010-07-22
Release 4.1	2010-04-27
Release 4.0	2009-12-21
Release 3.4 GA	2009-08-28 [Pillar Axiom 600 only]

1 Terms and Conditions of Use

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

2 Purpose

This document describes new features, capacities, configuration requirements, operating constraints, known issues and their workarounds, and other items for release 5.4 of Oracle Pillar Axiom 500 and 600 storage systems. The document covers hardware, firmware, software, cabling, and documentation. The information provided is accurate at the time of printing. Newer information may be available from your Pillar Axiom authorized representative.

3 Product Release Information^a

Release 5.4 is a special maintenance release of the Pillar Axiom software for the Pillar Axiom 500 and Pillar Axiom 600 SAN-only storage systems.

3.1 System Enhancements

This update provides enhancements and quality improvements to Pillar Axiom 500 and 600 SAN-only storage systems.

The overall system functions provided in this update are similar to those provided in release 5.3. In addition, this release provides improved system efficiency and robustness, as well as including a rollup of all defects fixed in customer patches up to and including release 5.4.

For the list of defects that this release resolves, see Table 11.

3.1.1 QoS Enhancements

This release provides an option that allows an administrator to balance QoS I/O contention across the Slammer ports. If this option is enabled (it is disabled by default) and a risk of I/O contention exists among multiple applications, the QoS service manager begins to prioritize application I/O requests sooner than in previous releases. When QoS management engages, applications with a High or Premium QoS setting will continue to operate at a higher service level than that for lower QoS priority settings.

In some cases, when a system is very heavily loaded, applications with a lower level of QoS might experience reduced service levels. This reduction in service level is to be expected with the QoS service manager delivering higher levels of service to High or Premium applications. When a system has low to moderate contention, the observed performance should be similar to that of systems running earlier software releases.

Note: If the service levels given to applications with lower QoS settings are insufficient, you can raise the QoS levels of the effected applications or disable the QoS balancing feature, which returns to the same I/O behavior as delivered by prior software releases.

Note: New Pillar Axiom customers are encouraged to enable this feature when the system is installed and enjoy the benefits of faster reacting, QoS I/O priority management.

This option is enabled and disabled by means of a Brick console command. To enable this feature, contact Oracle Customer Support. They will send you a Windows script that will enable or disable this feature.

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

3.1.2 Software Enhancements

- **Thin provisioning on SSDs**

Administrators can create logical volumes that have a maximum capacity of up to the maximum addressable logical capacity that is allowed for the OS, with little physical storage actually assigned to the volume. This characteristic is referred to as *thin provisioning*. As the Pillar Axiom system writes data to the thinly provisioned volume and capacity is consumed (called *in-fill*), the system automatically allocates additional capacity to the logical volume as needed.

Thin provisioning is now supported for the solid state drive (SSD) storage class. The thin provisioning feature has been and continues to be supported for all other storage classes.

- **Pilot management controller monitoring**

The Pillar Axiom system now monitors the status and reliability of the network connections of each Pilot control unit (CU). The Axiom system will detect a link failure on the active Pilot CU and after one minute will automatically fail over that CU to the passive CU. This feature supports continued management access to the Pilot without the need for manual intervention.

When the system detects a link failure on any of the management ports, the system will generate a system alert and reflect the problem in the GUI (see Section 3.7.1).

3.1.3 Hardware Enhancements

In the rare instance where both Pilot control units (CUs) fail, replacement of both Pilot CUs is now supported. Contact Oracle Customer Support for assistance.

3.2 Changes to Licensing

Beginning with release 5.0, all licensing requirements in Pillar Axiom systems, with the exception of the secure shell (SSH) license, have been removed. For more information on licensing, refer to Section 3.4.

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

3.3 Changes to How the Pillar Axiom Software Operates

3.3.1 SAN Replication

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

The release 4.x SAN replication product, which is named Pillar Axiom MaxRep Replication for SAN, is *not* compatible with the release 5.x AxiomONE Replication for SAN product. All customers who use SAN replication must utilize Oracle Advanced Customer Support Services to establish a new installation or to migrate from an existing replication solution.

3.3.2 Capacity Utilization Across Brick Types (Storage Classes)

In software releases prior to release 4.1, for a Pillar Axiom system containing a mix of SATA and Fibre Channel (FC) Bricks, individual logical volumes could utilize space from both the SATA and FC storage pools. As of release 4.1 and later releases, administrators cannot *create* volumes that span Brick types (called *Storage Classes* as of release 4.1). Legacy volumes created prior to release 4.1 that span Storage Classes, however, can continue to exist in the system after updating to release 5.4.

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

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

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

3.3.3 Reported Free Capacity

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

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

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

3.3.4 Method of Reporting Volume Sizes

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

- **Allocated logical capacity.** The amount of storage physically allocated for the logical volume. This value can change if the logical volume is migrated or copied with a different Quality of Service setting or if Thin Provisioning in-fill occurs. This value is the actual amount of storage allocated, minus RAID overhead.

- **Addressable logical capacity.** The size of the logical volume as visible to a client. This capacity defines the LBA range visible to a client and is guaranteed to remain constant across data migrations or volume copies that do not involve a volume resize operation.

3.3.5 Network TCP Port Configuration

After updating a Pillar Axiom system to release 5.4, administrators then access the updated system from a client workstation to download various utilities, including the new Java-based GUI, Pillar Axiom Storage Services Manager. (See Section 3.5.7.)

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

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

3.3.6 PDSCLI Is No Longer Supported

Beginning with release 5.0, the `pdsccli` Pillar Axiom command level interface is no longer available or supported. Instead, all scripting should use the new `axiomcli` interface.

3.4 Licensing Optional Premium Features

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

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

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

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

- Pillar Axiom MaxRep Asynchronous Replication - Integrated Software Option - Terabyte Metric (for factory installation)
- Pillar Axiom MaxRep Asynchronous Replication with Application Protection - Integrated Software Option - Terabyte Metric (for factory installation)
- Pillar Axiom MaxRep Synchronous Replication - Integrated Software Option - Terabyte Metric (for factory installation)
- Pillar Axiom MaxRep Synchronous Replication with Application Protection - Integrated Software Option - Terabyte Metric (for factory installation)

For additional information, please contact [Oracle Global Customer Support](http://www.oracle.com/support/contact.html) (<http://www.oracle.com/support/contact.html>).

3.5 Pillar Axiom Software Update

Pillar Axiom software updates are customer installable. If you prefer to have a technician from Oracle Advanced Customer Support (ACS) install the update, contact your ACS Field Sales Representative (see Table 5 Contact information).

The 5.4 release and patches can be installed by means of the GUI software update process or by means of the Axiom CLI (refer to the *Pillar Axiom CLI Reference Guide* for details).

If Oracle Customer Support has given you special installation instructions, follow those instructions using the information in this entire section as a guide.

Note: Updates to release 5.4.x should always be to the current recommended general release patch unless Oracle Customer Support has provided you a *specific* release or patch number.

3.5.1 Pillar Axiom GUI Supported Platforms

As of R5.x, the GUI client must be installed on a client workstation to administer the Axiom system. This client uses TCP port 26008 to communicate with the Pilot. Before you update an Axiom system from release 4.x to release 5.x, make sure all firewalls between your management workstations and the Axiom Pilot allow TCP port 26008 for all three of the Pilot IP addresses.

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

- CentOS Linux, version 5
- Linux Fedora Core
- Oracle Linux, version 5
- Red Hat Linux, versions 4 and 5
- Solaris 10, Sparc
- Solaris 10, 64-bit x86
- Ubuntu
- Windows Server 2003
- Windows Server 2008
- Windows 7
- Windows Vista
- Windows XP

Linux and Windows platforms require Java version 1.7.0 or higher.

Macintosh platforms require Java version 1.7.0_11.

Windows platforms require Java version 1.7.0 or higher. However, if you are using the MSI installer, Java doesn't need to be installed separately because the GUI, a stand-alone application, does not use Java as a browser plug-in but instead uses its own Java VM.

3.5.2 Preparing and Planning for the Update

When updating an Axiom system to release 5.4, the process for updating a system from release 4.x is not the same as the process for updating a system from a lower 5.x release. The following sections note those differences.

3.5.2.1 Updates From Release 4.x or Below

Important! The Axiom storage system must be running release 4.5.x or a higher 4.x release to update the system to release 5.4.x. For information on updating your system to release 4.5.x or higher, take one of the following actions:

- Refer to the *Pillar Axiom Customer Release Notes for Release 4.5*.
- Refer to the *Pillar Axiom Customer Release Notes for Release 4.6*.
- Contact Oracle Customer Support.

This update also requires a pre-update system audit. The Axiom storage system must be at release 04.00.02 (4.0.2) to perform this audit. After the audit is performed, do not change the Axiom storage configuration until after the software update. Contact Oracle Customer Support in advance of the update to perform the audit.

Important! All customers who use Pillar Axiom MaxRep for SAN on release 4.x must contact Oracle ACS to establish a new installation or migrate from an existing replication solution. (For more information, see Section 3.3.1.)

The update from release 4.x to release 5.x requires a maintenance period and can take up to 1.5 hours, depending on the system and the configuration size.

Only Pillar Axiom 500 and Pillar Axiom 600 SAN storage systems can be updated to release 5.4.x. The Pillar Axiom 300 and Axiom systems with NAS Slammers are not supported on release 5.x.

Updates to release 5.4.x might require updates to other applications or devices in your environment. Please refer to the current *Pillar Axiom Support and Interoperability Guide* or to the documentation for those applications or devices.

Note: For release 5.4.x, Pillar Axiom Path Manager on a client host does not need to be updated.

Software updates from release 4.x use the extracted “FN” package in the *2of3 package zip file (see Table 2 on page 13). Be sure that the first five characters of the name of the extracted file match your Axiom model number.

3.5.2.2 Updates From a Lower 5.x Release

Pillar Axiom systems consisting of three or four Slammers must perform a disruptive update.

Updates from qualified lower 5.x releases are non-disruptive and do not require an Axiom system maintenance period, *except* when any of the following characteristics is true:

- [13763238, 13765826, 13764221] If the installed release is below 05.00.05 (5.0.5) or is release 05.02.01 (5.2.1), perform the software update to 5.4.x disruptively.
- [14788300] If your Axiom system has three or more Slammers, the first time you update your system to 5.4.x, you must perform a disruptive update. Subsequent updates to a higher 5.4.x release can be performed non-disruptively.

Software updates on Axiom systems already at R5.x use the rpm package that you extract from the *10f3 package zip file. These update packages are not hardware specific; you can use the Axiom 600 or the Axiom 500 patch version.

3.5.2.3 Installed Software Module Versions

You can obtain your Axiom software version from the **Support > Software Modules** display in the GUI or from the `software_update -list` command in the CLI.

Note: The module names are slightly different in the CLI output. The only significant difference is that, for the firmware for the latest SATA Bricks, the GUI refers to *SATA V2* firmware whereas the CLI refers to *NEXTGEN* firmware.

Also, if your Axiom system is configured for Call Home, the software version is available to the Technology Support Center from the periodic Call Home logs.

In the GUI, the following Axiom software module names are displayed:

Table 1 Axiom software modules

Module name	Purpose	Versions
Pilot OS	The base operating system of the management controller	This module does not always change with every patch and does not need to match the Pilot or Slammer software.
Pilot software	The Axiom management applications	This module changes with every patch. The revision of this module must match the release or patch version. The revision of this module must match the Slammer software or the Axiom 600 Slammer software
Slammer software (AX500 or AX600)	The runtime protocol services software for the Slammer	This module changes with every release or patch version. The revision of this module must match the release or patch version and the Pilot Software.
Slammer PROM (AX500 or AX600)	The Slammer BIOS (Used to test and boot the Slammer CU.)	This version changes only when needed in a major release, such as R2.x, R3.x, R4.x, or R5.x. If this version changes, the software update must be done disruptively. When upgrading from release 4 to release 5, the PROM version will change for Axiom 500 Slammers but not for Axiom 600 Slammers. When upgrading from release 5.x to release 5.x, the PROM version will not change for either Slammer type.

Module name	Purpose	Versions
Brick firmware	The firmware on the RAID controllers	<p>The Brick firmware versions do not change with every release or patch. The exact name depends on the Brick type. If your Axiom system does not contain one of Brick types listed below, you might not see that Brick firmware component.</p> <p>The Brick can be any of the following types:</p> <ul style="list-style-type: none"> • Brick Firmware (R4) or Brick SATA Firmware (R5) • Brick SATA V2 Firmware (R4, R5) • Brick FC Firmware (R4) or Brick Fibre Channel Firmware (R5) • Brick FC V2 Firmware (R4) or Brick Fibre Channel V2 Firmware (R5) <p>Important! Only a single version of Brick firmware should be displayed for each Brick type. The last two fields of the firmware version should be identical for all Brick firmware modules.</p>
Drive firmware	The firmware on the drives	<p>The drive firmware version never changes with a release or patch. Drive firmware updates are done with special packages only.</p> <p>The special packages are in the *30£3 zip file for the patch or release.</p>

3.5.3 Obtaining Axiom Software

Axiom patches and releases are available on My Oracle Support (MOS) under the menu **Support > Patches & updates**. Use one the following URLs to access MOS:

- <http://www.oracle.com/support>
- <https://support.oracle.com>

You must have a valid MOS account with “Firmware” support to download Axiom software.

For assistance in locating Axiom patches and releases, search the Knowledge Base for the MOS document ID 1422199.1 and locate the article that is titled *Pillar Axiom: How to Find Pillar Axiom firmware and patches in MOS*.

Tip: If you have the MOS Patch Number for the software, enter it on the simple search screen.

The MOS Patch Numbers for the 5.4.0 release are as follows:

- Axiom 600 5.4.0 MOS Patch Number 16237327
- Axiom 500 5.4.0 MOS Patch Number 16245305

Axiom patches and releases can be classified as *General* or *Restricted*.

- General software can be downloaded by anyone with no password. However, if you have any questions regarding prerequisites, interoperability, functional changes, a system audit, or the MOS Patch Number, contact Oracle Customer Support.
- Restricted software requires a password to download. Contact Oracle Customer Support for a password.

In some circumstances, the software might be attached to your MOS SR, ready for downloading.

3.5.3.1 MOS Software Packages

Be sure you download the MOS package that matches your Axiom model when upgrading from release 4.x to release 5.x. Separate MOS patch numbers and packages are available for the Axiom 500 and Axiom 600.

Axiom software is available on MOS in three separate zip files with Release Notes or a patch readme.html file.

Important! *You only need one of those three zip files* (see Table 2). For the zip file that you select, you will (in a later step) extract the contents of that zip file to obtain the patch (or other software) and the release notes. In that later step, do not extract the tgz or rpm packages that are contained by that zip file.

The names of the zip files begin with the MOS patch number and the patch release version. The file names end with a string that identifies placement of the file in the sequence of the three zip files. For example, given a MOS patch number of, say, 16354312 and an Axiom release number of, say 5.4.11, you would see the following set of three files:

Table 2 Sample MOS software packages

Archive file name	Update type	Notations
p16354312_5411_Generic_1of3.zip	Release 5.x to release 5.4.x	Contains the Release Notes or patch readme.html file and the software in RPM format for updating the Axiom system. For example, the Axiom release 5.4.11 build 92 package would be in a file named AxiomONE-SW-050411-009200.i386.rpm. The Axiom model designation is not included in the RPM patch format. The RPM package type supports either the Axiom 500 system or the Axiom 600 system.

Archive file name	Update type	Notations
p16354312_5411_Generic_2of3.zip	Release 4.x to release 5.4.x	<p>Contains an Axiom 600 package or an Axiom 500 package in the release 4 model-specific package format for updating the Axiom system. Also contains the Release Notes or the patch readme.html file.</p> <p>Note: The release 4 to release 5 update packages in the zip file indicate the model number and the release. These package files are tgz files and must be staged as tgz files. For example:</p> <ul style="list-style-type: none"> • Ax500_FN_050411-009200.tgz For the Axiom 500 only and will update to release 05.04.11 build 92. • Ax600_FN_050411-009200.tgz For the Axiom 600 only and will update to release 05.04.11 build 92. <p>WARNING! Do not stage the wrong package when upgrading from release 4 to release 5.x. Doing so can render your Axiom system inoperable. Make sure the first five characters of the package name are the same as your Axiom model.</p>
p16354312_5411_Generic_3of3.zip	Special	<p>Contains special files such as drive firmware packages.</p> <p>Do not download the *3of3 package file unless instructed to do so by Oracle Customer Support.</p>

Important! Save the Release Notes or patch readme.html, because you will need to reference them when verifying the software staging and the update. The readme.html included in each of the three zip files is the same. That information is also the same as the information that is available on MOS.

3.5.3.2 Software Versions for the 05.04.00 Release

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

Table 3 Pillar Axiom software versions

Software module		Pillar Axiom model	Software version
Pilot OS		All	05.04.00
Pilot Software		All	05.04.00
Slammer PROM		500	03.00.00
		600	05.00.00
Slammer Software		All	05.04.00
Serial ATA (SATA)	SATA RAID controller, Version 1	All	07.23.13
Brick Firmware	SATA V2 RAID controller, Version 2 (includes SSD)	All	00.23.13
Fibre Channel (FC)	FC RAID controller, Version 1	All	01.23.13
Brick Firmware	FC V2 RAID controller, Version 2		02.23.13
Drive Firmware		<i>Version depends on drive vendor and model.</i>	

Note: Not all of the above information is necessarily displayed:

- If your Pillar Axiom system does not contain a particular Brick type (for example, Fibre Channel, SATA, SATA V2, or SSD), the GUI does not display the firmware entry for that Brick type.
- The GUI does not display drive firmware unless you are performing a drive firmware upgrade and have staged a drive firmware package.

3.5.3.3 Patch Readme Version Information

Each 5.4.x patch has a Readme file. This file is also included as the readme.html file in all three of the MOS zip files. The version information is in the section titled “Verifying Upgrade”.

For example, for the Axiom Release 5.4.11 build 92 patch, the version information would look like the following:

Software module	Installed version
-----	-----
Pilot OS	050411-009200
Pilot Software	050411-009200
Slammer PROM	030000-015001
Ax600 Slammer PROM	050000-040000
Slammer Software/Ax600 Slammer Software	050411-009200
Brick Firmware	012319-012319/002319-002319/023119-023119/072319-072319

The Brick firmware versions are listed in the following order:

- Brick Fibre Channel Firmware
- Brick SATA V2 Firmware
- Brick Fibre Channel V2 Firmware
- Brick SATA Firmware

3.5.4 Download the Update Package

1. From the My Oracle Support (MOS) web site, select the **Download** option.

You should see the three zip files.

2. Select the upgrade package:

Choices:

- The *1of3 zip file for upgrading from release 5.x to release 5.4.x.
- The *2of3 zip file for upgrading from release 4.x to release 5.4.x.

Important! When downloading the *2of3 zip file, make sure you have the correct model number.

3. Save the zip file.
4. Extract the readme.html file and the RPM or TGZ upgrade package from the zip file.

Important! Do not expand the tgz file in the *2of3 package. That file must be staged as the tgz file.

3.5.5 Stage the Update Packages

Set the Axiom session timeout to a value of at least 60 minutes during package staging.

Place the extracted rpm or tgz file in a convenient location on a workstation with a fast, reliable connection to the Axiom Pilot.

1. In the GUI, navigate to the **Support > Software Modules** screen.
2. Upload the update package.

Depending on your current release, choose one of the following actions:

- For R5.x systems, select **Upload Software Package**.
- For R4.x systems, select **Upload Package**.

3. Navigate to the location of the software.

Make sure you are uploading an .rpm file to R5.x systems or a .tgz file to R4.x systems.

Important! The first five characters of the tgz file on R4.x systems *must match* the Axiom model number. Do not stage a file that begins with "Ax600" on an Axiom 500 system. Do not stage a file that begins with "Ax500" on an Axiom 600 system. If you accidentally stage the wrong file type, contact Oracle Customer Support for a special file named `AxX00_N9_999999-999999.tgz`, which must be staged to recover *before* you stage the correct package file.

4. After the package staging completes, check the Axiom system to make sure the system did not generate an alert indicating that package staging failed.

5. Go to the **Support > Software Modules** screen and verify staging.
 - On release 5.x systems, the staged software is displayed below the installed software.
 - On release 4.x systems, you must select **Update Software** to see the new staged software. Make sure you click **Cancel** when leaving this screen to avoid starting the update process.
 - The Axiom CLI can be used to display the contents of the /STAGED software.

Make sure that the staged software versions are the same as those in these Release Notes or the Patch readme.html file.

Every patch contains all software components. You should see *all* of these components in the staged software. If any component is missing, contact the Oracle Customer Support Center and do not initiate the upgrade.

- Pilot OS. Must match the Release Notes or Patch readme.html.
- Pilot Software. *Must* match the patch version and the Slammer software.
- Slammer Software (Axiom 500 systems). Must match the patch and Pilot Software version.
- AX600 Slammer software (Axiom 600 systems). Must match the patch and Pilot Software version.
- AX600 Slammer PROM. *Must* always be 05.00.00-0400.00. Should not appear on release 4.x to release 5.x updates of Axiom 500 systems.
- Slammer PROM (Ax500). *Must* always be 030000-015001.
- Brick Firmware. You must see a Brick Firmware version for every type of Brick installed on your Axiom system. The last two fields must match on all Brick types.
- Brick Disk Drive firmware versions. Displayed only on release 4.x systems or on release 5.x systems when you are installing a drive firmware upgrade with the assistance of Oracle Customer Support.

For every hardware component in your Axiom system, make sure that the “Applies to Current Hardware” flag indicates Yes.

If you are not initiating the update at this time, exit the screen.

3.5.6 Install the Update

If you are performing a disruptive update, prepare your clients for loss of access to data.

1. Shut down all applications and clients.
2. Disconnect the external SAN interfaces, or block access at the SAN switches.
3. Wait at least 15 minutes before you begin a disruptive update.

This wait will allow all LUNs to be placed in an idle state. (Sometimes the status is shown as Conservative.)

Before you initiate the upgrade, check the status of all hardware components and data.

1. The status of the Axiom system must be Normal with all components in a Normal state.
2. The system status icon in the bottom left of the GUI must be green.

3. Check for any system notifications.
 - For release 4, check for Administrator Actions.
 - For release 5, check for Alerts.

If an Administrator Action or Alert is present, a flashing amber triangle will appear at the bottom left of the GUI. The Axiom system will not begin an update if any component has a Warning or Critical status or if any Administrator Actions or Alerts exist.

Before you initiate the upgrade, make sure that the Staged Software versions match the 5.4.0 Release Notes or the 5.4.x Patch readme.html. See the above on Staging Software.

On release 4.x systems, disable the Call Home feature to prevent a delay of the update. If you want to preserve any logs or events, collect a manual log set and save it before you begin the update.

Note: When updating a system from release 4.x to release 5.x, all events and logs will be cleared. You will not be able to access release 4.x events or logs after the Axiom Pilot has been upgraded to release 5.x.

If you are upgrading from release 4.x, the Axiom GUI should automatically select all of the software modules and the **Restart System** option. All of the following modules and options must be selected:

- Pilot OS.
- Slammer Software Ax600 on the Axiom 600 only.
- Slammer Software on the Axiom 500 only.
- Slammer PROM AX600 on the Axiom 600 only.
- All Brick Firmware modules for your installed hardware.
- The Slammer PROM should *not* be selected on the Axiom 500. It does not change on this update.
- **Restart System** must be selected.

Important! Do not unselect any of the above modules or options; otherwise, you might render your Axiom system inoperative. Do not exit this GUI screen as on 4.x releases below 4.6.8; otherwise, doing so might result in changing these selections.

Click **OK** to start the update.

If you are updating from release 5.x, navigate to the **Support > Software Modules** screen and select **Actions > Update Software**.

If you are logged in as a support account, you will see an **Install Action** column. Leave this column set to **Install if newer version** unless you are instructed by Oracle Customer Support to change this to **Do not Install** or **Always Install**.

Inspect the **Installed** and **Staged Version** columns. If they differ, the Axiom system should automatically install the higher version.

During the installation, you should see the following changes:

- The Pilot OS changes on an initial update to release 5.4.x. After the initial update, the Pilot OS will only change if noted in the patch readme.html.
- The Pilot Software and Slammer Software or Ax600 Slammer software always changes on any release or patch installation. The staged version of these modules must be identical.
- The Slammer PROM or Ax600 Slammer PROM never changes on any release 5.x to release 5.x update.
- The Brick Firmware versions might or might not change. Verify the Staged Version against the Release Notes or Patch readme.html information.
- Disk Drive firmware should never be displayed unless you are doing a disruptive drive-firmware update. If you see drive firmware displayed, cancel the release update and stage the correct rpm software package from the *1of3 zip file.

To perform a *non-disruptive* update, leave all modules in their default status. Make sure the **Update software without restarting system** option is selected. Do not select any of the Software Update Options unless instructed to do so by Oracle Customer Support.

To perform a *disruptive* update, select the **Restart and update software** option. Also, select the **Shutdown Slammer** and **Ignore current requests** options. Do not select any other options unless instructed to do so by Oracle Customer Support.

After the Pilot connection becomes available during the update, you need to download and install the 5.4.x GUI client. To download the 5.4.x GUI client, use the `http:// AxiomIP_or _Hostname` URL. This URL identifies the public IP address or the name of your Axiom system.

Note: Additional features in release 5.4.x might hang the 5.3 and older GUI clients if you attempt to view the status of the Pilot.

The Axiom system will prepare the system for the update and then begin the update. On release 5.x to release 5.x updates, you will see a *prepare for upgrade* task in the task bar. While the completion status of that task is below 30%, the Axiom system is waiting for a task or operation to complete before it will initiate the update. If a release 4.x system is in the process of sending a Call Home notification, the system will not initiate the update.

A disruptive update shuts down the Slammers before the update begins. This step is mandatory for a release 4 to release 5 update to avoid Slammer CUs becoming disabled due to a change in the memory allocation. A non-disruptive update leaves the Slammers running to provide data access. During the update, the new software is placed in the Slammer memory, and the Slammer CUs are warm started to enable the new software.

The first component that is updated is either the Pilot Software or the Pilot OS, or both. This action allows any fixes in the update process to be effective in the higher release. After the Pilot restarts using the new Pilot Software, the Pilot initiates the rest of the update, with the component order depending on the update type.

As the Slammers are updated, the Pilot inhibits any Slammer from restarting until the Pilot is ready to process their startup sequences. If a Slammer CU experiences a software fault during a non-disruptive update, the CU is not allowed to recover until the rest of the update completes.

If the update is disruptive, every RAID CU is shut down and all firmware updated at the same time. Then the RAID CUs are restarted.

If the update is non-disruptive, all of the Brick RAID CU0s are updated and restarted, followed by all of the Brick RAID CU1s. This action sequence allows data access to continue during the update.

The status of the data and the system components, as displayed in the GUI, might not be accurate until the Axiom system completes the final Boot State Pilot at the end of the update. The task list can be used to monitor the update if you have a release 5 client installed.

Tip: Using `http://AxiomIP_or_Hostname`, you can monitor the release 5.x Axiom system with a web browser, including the browser on a smart phone. Using this HTTP connection, you cannot modify the configuration of the Pillar Axiom system.

3.5.7 Install the GUI Client

After the update completes, download and install the 5.4.x GUI from the Axiom system using `http://AxiomIP_or_Hostname`.

If you have a previous GUI client installed, you will receive a warning that your client is incompatible with the installed software. You can override this warning and continue, but the display might not be accurate. In some cases, you might receive an error, in which case you will need to install the latest GUI client.

Note: Any attempt to view the Pilot status might hang your client due to new features in the 5.4 Pilot Status display.

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

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

3.5.8 Verify the Update

1. Enable Call Home, if it was disabled.
2. Verify that the Call Home settings have not changed.
3. Check the Axiom hardware and software status.

The System Status on the **Configure > System** screen should be Normal. The system status icon on the bottom left side of the GUI should be Green. Make sure that the amber triangle at the bottom of the GUI is not flashing, which would indicate that an Administrator Alert is present.

4. Check the status of your LUNs on the **Configure > LUNs** screen.

All LUNs should be Online. No LUN should display an Offline or Conservative status.

5. Check your Software Module versions against the 5.4.0 Release Notes or the Patch readme.html file.
 - You should have only one version of Brick firmware for each brick type. The final two fields for all Brick Firmware types should be the same.
 - Your Slammer Software or AX600 Slammer Software version should match the Pilot Software.
 - The Pilot OS version should match the Release Notes or Patch readme.html
 - The Slammer PROM should be 030000-015001 for an Axiom 500 system and the AX600 Slammer PROM should be 050000-040000.
6. Verify data access from your applications and clients.

3.5.9 Upgrade Recovery Using Override Software Update Options

The “Software Update Options” can be used to recover from simple update failures.

3.5.9.1 Failed Brick Firmware Update

If the Axiom fails in updating the Brick Firmware components, you might see multiple Brick Firmware versions displayed for the same Brick type. The ability to run multiple versions is tested before any Software release. As long as the Axiom system returns to a status of Normal, with all components online, you can safely re-run the upgrade.

1. Remove any Administrator Alerts.
2. Select **Update Software without restarting system**.
3. Select **Override failed software update**.
4. Restart the update.

Important! Before attempting to recover from any other update failure, contact Oracle Customer Support for your recovery options.

The available software update options are as follows:

- **Restart and update software**
This option forces a disruptive upgrade.
- **Ignore Compatibility**
Use this option only as instructed by Oracle Customer Support.
- **Shutdown Slammer**
This option will forcibly shut down the Slammers and is typically used only when a hung process in a Slammer CU is preventing normal shutdown. This option should only be used as instructed by Oracle Customer Support, as it can cause pinned data or cause the Lost Data flag to be set.
- **Ignore hardware status**
You must have two functional Pilot CUs for any upgrade. You can ignore the hardware status of any other component so that you can install a potential fix that will help recover that hardware. For example, Oracle Customer Support might direct you to select this option so you can install a patch to prevent repeated Slammer CU software faults or to allow a RAID CU to recover. This option should only be used as instructed by Oracle Customer Support.

- **Ignore system alerts**
The Axiom system will not initiate an update if an Administrator Alert exists. Contact Oracle Customer Support with the details of any Administrator Alert before using this option.
- **Ignore current requests**
This option terminates any commands or tasks that are in progress in the Pilot software and launches the update. Contact Oracle Customer Support before using this option.
- **Override failed software update**
This option clears the software update status. Except as noted in retrying an update due to a failed RAID firmware update in which all components return to Normal status, contact Oracle Customer Support before using this option.

If you are instructed to set any Software Update Option, the Axiom system will display **Software Update Validation Errors** below the set of options when you select incompatible options. The reason for the validation error is displayed. The Axiom system does not initiate a software update if there are any software update validation errors. Contact Oracle Customer Support with the validation error information for assistance.

3.6 Pillar Axiom CLI Supported Platforms

As of release 5.4, the Pillar Axiom CLI (`axiomcli`) is supported on the following platforms:

- AIX 5.3 (TL 7 and later), 64-bit PowerPC
- AIX 6.1 (TL 3 and later), Power
- Citrix 5.6 XenServer, x86_64
- HP-UX 11i, ia64
- Mac OS X 10.6, x86_64
- Open SuSE 11.3, x86
- Red Hat, CentOS, Oracle Linux, versions 4.8, 5.4, and 6.x, x86
- SLES 11, x86
- Solaris 9, Sparc
- Solaris 10, Sparc
- Solaris 10, x86
- Solaris 11, Sparc
- Solaris 11, x86
- Windows 2003, x86
- Windows 2008 R2, x86
- Windows 7, x86
- Windows 8, x86

- Windows Vista, x86
- Windows XP, x86

You can download management software for the Pillar Axiom system, such as the CLI, as follows:

1. Point your browser to [http:// AxiomIP_or_Hostname](http://AxiomIP_or_Hostname), where *AxiomIP_or_Hostname* is the public IP address or the name of your Pillar Axiom system.
2. Click **Management Software**.
3. In the Axiom Command Line Interface (CLI) section, click the appropriate download link.

3.7 New Features in the Pillar Axiom GUI

3.7.1 Pilot Network Connectivity

The new **Monitor > Pilots > Actions > View Pilot** menu action provides information concerning network connectivity for a selected Pilot control unit (CU). This information indicates the status of the Ethernet, management, serial, and service ports for that Pilot CU.

When a network connection for a Pilot CU is lost, the Pillar Axiom system generates an alert and a system event. For example, if a serial port (used for serviceability) or a private management interface (PMI) Ethernet port fails, the system generates an appropriate system alert so an administrator can correct the lost connection.

3.7.2 Slammer Port Connectivity

[15993210] On the **Host - LUN Mapping** and **LUN - Host Mapping** screens, the columns displaying Slammer port connectivity information have been removed. The port columns have also been removed from the table on the Mapping tab of the **Create LUN** and **Modify LUN** dialogs.

The port connectivity information is now available through the **Host - LUN Mapping** and **LUN - Host Mapping** panes by clicking **Actions > Initiator-to-Slammer Connectivity**.

3.7.3 Clone Repository Capacity

On the Data Protection tab in the screens for creating, modifying, and cloning LUNs, the default addressable capacity for a Clone LUN has been changed from 10 GB to 100% of the addressable capacity of the LUN. If the administrator overrides this default, the Axiom system stops automatic updates to the repository capacity when the user changes the maximum addressable capacity for the LUN. If the user then changes the addressable capacity for the LUN back to the default, the system resumes the automatic updates.

3.7.4 Deletion Dialogs for Deleting LUNs and Clone LUNs

The confirmation dialogs for deleting LUNs and Clone LUNs have been modified to allow for the deletion of multiple objects.

3.8 Pillar Axiom Path Manager Software

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

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

Note: APM requires TCP port 26004 to be open between the SAN client host and the Pilot management controller.

The current release of APM for SAN hosts varies by platform, as shown in Table 4.

Table 4 APM support

Operating system	OS release	APM release	Platforms
AIX	5.3 6.1 7.1	3.1	All platforms
Community Enterprise Operating System (CentOS)	5.6	3.4	32-bit x86, 64-bit x86
	5.7	3.5	32-bit x86, 64-bit x86
	5.8	3.6	32-bit x86, 64-bit x86
	5.9	In development	32-bit x86, 64-bit x86
	6.1	3.1	32-bit x86, 64-bit x86
	6.2	3.2	32-bit x86, 64-bit x86
	6.3	3.3	32-bit x86, 64-bit x86
HP-UX	11i v3	In development	All platforms
Oracle Linux (OL)	5.6	3.4	32-bit x86, 64-bit x86
	5.7	3.5	32-bit x86, 64-bit x86
	5.8	3.6	32-bit x86, 64-bit x86
	5.9	In development	32-bit x86, 64-bit x86
	6.1	3.1	32-bit x86, 64-bit x86
	6.2	3.2	32-bit x86, 64-bit x86
	6.3	3.3	32-bit x86, 64-bit x86

Oracle Solaris	9	2.1	SPARC
	10	3.0	All platforms
	11		
Oracle VM Server for x86	2.2	3.2	All platforms
	3.0	3.2	All platforms
	3.1	3.3	All platforms
Red Hat Enterprise Linux (RHEL)	5.6	3.4	32-bit x86, 64-bit x86
	5.7	3.5	32-bit x86, 64-bit x86
	5.8	3.6	32-bit x86, 64-bit x86
	5.9	In development	32-bit x86, 64-bit x86
	6.1	3.1	32-bit x86, 64-bit x86
	6.2	3.2	32-bit x86, 64-bit x86
	6.3	3.3	32-bit x86, 64-bit x86
SUSE Linux Enterprise Server (SLES)	11 SP2	3.1	32-bit x86, 64-bit x86
Windows	Server 2003 Server 2003 R2 Server 2008 Server 2008 R2	3.3	All platforms
	Server 2012	In development	All platforms

Note: OS releases earlier than those listed above are not supported. OS releases later than those listed above will be supported in future APM releases.

Tip: For optimal performance with the 5.4 release of the Pillar Axiom software, Pillar recommends that APM clients use round-robin access from the host. Doing so ensures that all ports on the Slammer CPU on which the LUN resides are fully utilized.

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

For release information, refer to the *Pillar Axiom Path Manager Installation Guide and Release Notes* for your platform, which can be found on the [Oracle Technical Network](http://www.oracle.com/pls/topic/lookup?ctx=pillardocs) (<http://www.oracle.com/pls/topic/lookup?ctx=pillardocs>) in the Oracle SAN Storage Systems section.

For the latest information on supported platforms and hardware, you can contact your Oracle Customer Representative. Also, you can access the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Oracle Technical Network](http://www.oracle.com/pls/topic/lookup?ctx=pillardocs) (<http://www.oracle.com/pls/topic/lookup?ctx=pillardocs>) in the Oracle SAN Storage Systems section.

4 Support

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

Table 5 Contact information

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

4.1 Supported Hardware Components in a Pillar Axiom System

Only Oracle-supplied parts for Pillar Axiom systems are supported. Hardware that does not conform to Pillar Axiom specifications or is not an Oracle-supplied part voids the warranty and might compromise data integrity.

4.2 Access to Pillar Axiom Systems

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

- The Pillar Axiom Storage Services Manager (GUI)
- The Pillar Axiom Command Line Interface (CLI)
- A plug-in or interface from a storage management client (SMI-S, Oracle ASM, Microsoft VSS, and so on)

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

4.3 Configuration Documentation

For information on the connectivity and interoperability of Pillar Axiom systems with various third-party software and hardware, contact your Oracle Customer Representative. Also, you can access the *Pillar Axiom Support and Interoperability Guide*, which can be found on the [Oracle Technical Network](http://www.oracle.com/pls/topic/lookup?ctx=pillardocs) (<http://www.oracle.com/pls/topic/lookup?ctx=pillardocs>) in the Oracle SAN Storage Systems section.

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

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

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

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

- In the Pillar Axiom Storage Services Manager (GUI), navigate to **Support > Documentation**.
- Point your browser to <http://system-name-IP>, where *system-name-IP* is the name or the public IP address of your Axiom system. Then, click **Documentation**.
- Log in to the [Oracle Technical Network](http://www.oracle.com/pls/topic/lookup?ctx=pillardocs).
<http://www.oracle.com/pls/topic/lookup?ctx=pillardocs>

5 Pillar Axiom System Limits

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

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

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

Consult with Oracle ACS to plan your Pillar Axiom system configuration prior to actual installation and configuration.

5.1 Pillar Axiom System Operating Limits

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

Table 6 Operating limits for all Pillar Axiom systems

Item	Description and range
Volume groups	Minimum = 1 Maximum = <ul style="list-style-type: none"> • 5000 total, out to five levels • 100 subgroups for a given volume group
Repository VLUNs ^b	Maximum = 1024
VLUNs	Maximum = <ul style="list-style-type: none"> • 8191 for each Slammer • 8191 for each system
Storage Domains	Maximum: 64 for each system
Number of Bricks in a Storage Domain	Minimum = <ul style="list-style-type: none"> • Primary domain: 2 Bricks^c • All other domains: 1 Brick Maximum = <ul style="list-style-type: none"> • SATA Bricks: 64 • FC or SSD Bricks: 32

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

^c For single-Brick systems, the primary Storage Domain consists of one Brick

Table 7 Operating limits for Pillar Axiom SAN systems (continued)

Item	Description and range
SAN LUNs	Maximum = <ul style="list-style-type: none"> • 8191 visible for each system • 256 visible for each host • 8191 visible for each SAN Slammer
SAN LUN size	Minimum = 1 to 2 GB. The exact value depends on these factors <ul style="list-style-type: none"> • Brick type (Fibre Channel, SATA, or SSD) • RAID geometry (RAID 5 or Pooled RAID 10) • Strip size (1 MB or normal) Maximum = system capacity (unlimited when using Thin Provisioning)
Volume Copies (full block snapshots)	Maximum = number of unallocated SAN LUNs, up to maximum SAN LUNs <ul style="list-style-type: none"> • 1024 for each system • 12 active for each LUN
Clone LUNs (partial block snapshots)	Maximum = number of unallocated SAN LUNs, up to maximum SAN LUNs
Clone repositories	Maximum = ¼ of the maximum number of SAN LUNs
iSCSI	Maximum = <ul style="list-style-type: none"> • 256 TCP connections for each iSCSI port • 256 iSCSI initiators for each iSCSI port • 2048 iSCSI initiators for each Slammer • 256 LUNs for each initiator • 32 persistent reservation registration keys for each LUN • 512 simultaneous commands for each iSCSI port
LUN mappings	Maximum = 512K for each Pillar Axiom system

5.2 Slammer and Brick Configuration Limits

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

Table 8 Brick configuration limits

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

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

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

Note: The maximum number of Bricks in a Brick string, depends on the Storage Class:

- For SSD Bricks, the maximum is 2.
- For all other classes, the maximum is 8.

For Fibre Channel (FC) Bricks:

- Version 1 FC Bricks are available as FC RAID Bricks and FC Expansion Bricks. Each Version 1 FC RAID Brick supports the attachment of a single Version 1 Expansion Brick.
Note: Version 2 FC Bricks do not have an FC Expansion Bricks option.
- A given Brick string can contain up to four FC Bricks (RAID or Expansion). Maximum number of FC Expansion Bricks in any string, however, is two.

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

6 System Requirements

6.1 Windows Requirements

6.1.1 Using Browsers on Windows XP Operating Systems

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

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

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

6.1.2 Memory Requirements

[13755138] For the Pillar Axiom Storage Services Manager graphical user interface to operate correctly, ensure that the client host has a minimum of 2 GB memory.

6.2 Network Requirements

6.2.1 Pilot Network Requirements

The Pilot management controller requires:

- Two 100 BaseT or two 1 Gb 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.

Note: All new system installations have the Enhanced Pilot, which supports 1 Gb connections. To find out whether an existing Pilot supports 1 Gb, contact Oracle Customer Support.

- Three IP addresses on the same subnet: one IP for each physical interface and one shared IP.

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

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

All ports used for incoming and outgoing communications with the Pilot management controller are described in Table 9.

Table 9 TCP ports used by the Pilot

This port number ...	Is used by this application ...	Notations
22	Call Home	Used by the standard Call Home option. If used, must be open to the Pillar Call Home server. (Compare port 443.)
25	SMTP	For email notifications sent by the Pillar Axiom system.
53	DNS	Used, for example, to resolve the host name of the Call Home server.
80	HTTP	Used for status monitoring and download of documentation and various applications, such as the Pillar Axiom Storage Services Manager (GUI).
123	NTP	Used to synchronize the time on the Axiom system with an external time service provider.
161	SNMP	Used by a host client to monitor the Axiom system.
162	SNMP traps	Used to transmit SNMP traps to an external monitor. Used by compliant UPS devices to warn the Axiom system that the UPS is running on battery.
427	SMI-S, UDP, and TCP service locator	
443	HTTPS	Used by the HTTPS Call Home option. If used, must be open to the Call Home server. (Compare port 22.)
5988	SMI-S clients	
5989	SMI-S clients	
26004	APM clients	
26008	Axiom CLI clients GUI clients	

6.2.2 Slammer Network Requirements

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

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

6.3 Power-Off Requirements

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

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

6.4 Power Cycling

Contact Oracle Customer Support before turning off and then turning back on (power cycling) a Pillar Axiom system except in the event of an emergency. In an emergency, turn off all power and then contact the Support Center. Contact the Support Center before touching any power cables or switches. There are some situations where *not* power cycling the entire system is the correct action.

For *failure* testing, do not power cycle individual components without first contacting Oracle Customer Support.

See also Sections 9.45, 9.50, 9.59, and 9.66.

7 Known Issues

The Pillar Axiom server issues listed in Table 10 are known at the time of this release. They are planned for resolution in upcoming releases. When available, Oracle will provide updated software or hardware.

For additional information or help on any of the issues below, please contact your Oracle Pillar authorized representative (see Table 5 Contact information).

Table 10 Known Pillar Axiom server issues

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
All	[13754238] The Pillar Axiom CLI product has no option to send a request and get the response status at a later time. All requests are synchronous.	This issue will be fixed in a future release.
	[13753767] The Pillar Axiom CLI product has no option to return the list of valid error codes.	This issue will be fixed in a future release.
	A series of Slammer CU warm starts on larger systems may result in the system going to a conservative status, where the system configuration becomes read only.	Contact Oracle Customer Support for assistance. If the system configuration becomes read only, the Axiom must be shut down and restarted. This issue will be fixed in a future release
	[16045865] Attempts to view the SAN Hosts and Host Groups with management applications that use the Axiom WebCLI will fail	This issue will be fixed in a future release.
	[14667547] The help information for the following Axiomcli –modify requests incorrectly shows –name instead of –newName: account, brick, event_notification, host_group, job, san_host, snmp_host, system, and ups	This issue will be fixed in a future release
SAN	[13762901] The configuration of the Pillar Axiom MaxRep Replication for SAN engine is closely coupled with the World Wide Names (WWNs) of the FC HBA ports. This means that, if the HBAs are replaced, the configuration must be redone manually.	If one must replace the FC HBA, reconfigure the replication engine. This issue will be fixed in a future release.

Slammer type	Pillar Axiom server issue or impact	Workaround or planned fix
	[14594661] If a copy of a LUN is created with a named protection schedule and the copy is deleted, the schedule cannot be deleted. Attempts to create another protection schedule with the same name will fail.	This issue will be fixed in a future release
	[13762697] There is no assurance that the labeled physical Ethernet ports on the Pillar Axiom MaxRep Replication for SAN appliance will match the software configuration of the ports.	Log in to the appliance and check the port configuration at the OS level to get the needed information. This issue will be fixed in a future release.
	[16059365] Attempting to create a protection schedule for a clone of a clone will fail.	This issue will be fixed in a future release.
	[15965242] Under some conditions, the Axiom Slammer may not complete a warmstart in time to resume normal SAN traffic.	This issue will be fixed in a future release
iSCSI	[13746664] On rare occasions, the iSCSI network interface card can hang, which causes the system to warm start.	Replace the iSCSI card, or contact Oracle Customer Support for assistance. This issue will be fixed in a future release.

8 Resolved Issues

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

Tip: For a complete listing of resolved issues, please contact Oracle Customer Support (see Table 5 Contact information).

Table 11 Resolved Pillar Axiom server issues

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
5.4	SAN	[16171956] A RAID Controller may place a drive in Temp Degraded Mode and never exit that state if additional drive recovery is required while in the temporary degraded mode.
		[16167200] The following statement in the Administrator Help for Thinly Provisioned Volumes no longer applies: Note: Solid-state drive (SSD) Bricks do not support thinly provisioned volumes.
		[16021839] RAID Controllers may experience software faults on RAID10 resources when there are write commands with very large block counts.
		[16019313] A race condition in recovering from drive faults when a preemptive CopyAway operation is initiated on a marginal drive may result in the rebuild to the spare drive not completing successfully.
		[16018397] Adds support for a new model of SSD drive.
		[16001466] Adds rotating patterns to the traffic generator for the Pitman SSF diagnostic utility.
		[15985123] A diagnostic mode has been added to the Pitman utility to increase internal traffic levels to help isolate faulty components in the SSF Private Interconnect.
		[15993210] The Axiom GUI client may hang after consuming nearly 1 Gigabyte of memory when viewing LUN mappings on systems with a large number of LUN mappings. .The change removes the Slammer port information from the Host to LUN, the Mapping tab on the LUN Create and Modify, and LUN to Host mapping screens. .Slammer port connection information has been moved to the "Initiator to Slammer Connectivity" sub menu on the Host and LUN Mapping screens.
[15980201] Fibre Channel Drives which drop their FC connection momentarily and return will be intentionally faulted if their sense code information on returning is invalid.		

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		[15952822] Creation of LUNs using the VDS Provider for Microsoft Virtual Disk Service fails.
		[15947014] The Axiom CLI -verify option always returns an "Ambiguous command name" and "Command Failed" status.
		[15923499] After a Pilot CU failover, events may not be displayed in chronological order in the GUI or in the response to an Axiom CLI request.
		[15922476] RAID controllers may experience intermittent software faults during maintenance operations where a data drive has been replaced with a drive of a larger size.
		[15904119] Pillar Axiom Storage Services Manager clients may not be able to access the Axiom system on client operating systems that support both IPV4 and IPV6 if Java 7 has been installed.
		[15899521] Slammer CUs may warmstart when processing events related to Storage Enclosure (Brick) Stuck Busy detection.
		[15887344] RAID Controllers may experience a software fault if a rebuild to the spare drive due to a drive fault overlaps with handling a temporary error condition on another drive in the same RAID array.
		[15880408] If an Axiom system has been upgraded from a 4.x release lower than 4.5 or a 5.x release lower than 5.2, LUNs created with a QoS of Random Write, where a RAID10 License was not installed at the time of the creation, may display incorrect values of 50% for LUN Overhead and Mirroring for Disk Protection rather than the actual values of 9 to 17% and RAID5 Parity. The LUN can be migrated using the GUI if RAID10 is desired for the Random Write QoS.
		[15868342] The Allocated Capacity displayed in the GUI and the CLI for LUN clones has been removed from the GUI and the CLI because the value was only that for the metadata used to manage the clones. The Used Repository Capacity should be used to manage clone space.
		[14835654] The Axiom CLI lun -add command has been modified to automatically configure cloneCapacity equal to the addressable capacity of the LUN if a value for clone capacity is not included. If a value is specified, including a value of 0, then that value is used.
		[14804469] Pilot CU Replacement fails. The replacement pilot is not updated to match the Pilot OS or Pilot Network configuration of the surviving Pilot.
		[14788300] The Slammer CU-to-CU SSF connections on Axiom systems with 3 or 4 Slammers does not properly enable and load balance the available links. A Disruptive upgrade to 5.4.x or a shutdown and restart after the upgrade is required to install the fix for this issue.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		[14783368] Slammer CUs may warm start if duplicate command requests are detected on the internal SSF CU to CU paths. The warmstart is replaced with a temporary path soft failure to recover.
		[14781343] A race condition in the Storage Enclosure RAID Firmware may result in multiple RAID Controller software faults during the processing of a stuck busy condition.
		[14757898] SAN Slammer CUs may warm start if an iSCSI client attempts to re-establish a connection and the attempt fails but a subsequent attempt succeeds.
		[14757689] Slammer CUs may warm start or Pilot CUs may fail over due to an inconsistent connection state on the internal private management network.
		[14757687] During a preemptive Copy Away operation, if the drive being copied to the spare continues to generate drive errors, both RAID controllers on the Storage Enclosure may experience repeated software faults.
		[14732091] If a Storage Enclosure (Brick) is removed from the system configuration, the management software may continue to poll every 5 minutes for Brick Port Status on the removed Enclosure until a Pilot CU failover occurs
		[14705899] If write commands are issued on the same range of a LUN in rapid succession, a read command to that same range may return cached data instead of the most recent write content.
		[14698062] The Axiom CLI clone_lun command incorrectly required all lower case for the -accessBias and -ioBias options.
		[14679162] If a Pilot CU fails over to the standby CU, the statistics information that has already been collected is lost.
		[14668450] Attempts to delete Clone LUNs with the Storage Services Manager may fail if the parent of the clone is mapped to a host.
		[14667652] Deleting a large number of resource objects such as Hosts, LUNs, or Clones can cause the System Configuration Database to become unresponsive, causing all further tasks that attempt to modify the system configuration to hang. The Pilot may become hung in Boot State Pilot if there is a Pilot CU restart.
		[14660847] The Temporary Degraded Mode (TDM) drive recovery method has been expanded to address additional data cache flush issues.
		[14641602] A race condition in handling Slammer CU Failover/Failback may result in additional Slammer CU software faults on the surviving Slammer CUs.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		[14637281] The Axiom system may allow staging of an improper rpm package and upgrade attempt. Subsequent recovery attempts at staging the correct package and selecting all upgrade overrides will fail due to incomplete cleanup from the improper staging.
		[14612252] The SMI-S provider TCP ports 5988 and 5989 may not be properly enabled in the Pilot internal firewall, causing loss of access for SMI clients or the Enterprise Manager Plug-In.
		[14610847] The Axiom may report Port Group Failovers on newly added Bricks as they are configured for use if the RAID Firmware on the Brick was updated automatically during the installation because the Firmware update procedure does not initialize the Port Group access after the Firmware update.
		[14597276] Storage Enclosure RAID Controllers may experience repeated software faults due to a lock contention issue in attempting to recover from a Stuck Busy condition.
		[14593620] The IBM SAN Volume Controller may lose path connections to the Axiom system after a Slammer CU warmstart.
		[14581579] A SAN host with multiple initiators and a large number of mapped LUNs may not be able to complete an Axiom Path Manager login, leaving the host in Not Communicating status.
		[14562156] The PITMAN SSF diagnostic utility may only generate traffic on a single internal Fibre Channel path when run in the disruptive maxtraffic mode.
		[14559815] Commands between the Slammer and Brick may be retried repeatedly, preventing normal data access.
		[14546662] Software upgrades may fail due to failure to drain the operation pool.
		[14540909] If a Pilot CU fails over, existing events may not be recovered properly, resulting in duplicate or lost events.
		[14539541] Deleting SAN LUN Clones may cause SAN Slammer CUs to experience repeated software faults.
		[14538971] A Brick may display all components as Critical status after replacement of the Brick SES module.
		[14538879] After an upgrade from Release 4.x the Call Home function might stop working due to dropping the trailing <i>m</i> from the Call Home server name of <code>callhome.support.pillardata.com</code> , resulting in <code>callhome.support.pillardata.co</code>

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		[14526786] The Statistics Tool might fail to parse and format statistics bundles due to missing or incomplete data samples.
		[14490781] Deleting a SAN Clone LUN with active host I/O to that LUN may cause host operating systems to continually enter an Auto Contingent Allegiance condition and retry of the I/O.
		[14489563] The Axiom system generates Brick Strings Unbalanced events once every 24 hours if the number of Bricks in any string differs by two or more from the number of Bricks in any other string. This Event is suppressed in this release.
		[14463206] Pilot CU network connection status and reliability of connection information has been added to the GUI and the axiomcli pilot -list -details response.
		[14405199] On large systems with a large number of LUNs and enough I/O to cause busy conditions on Bricks, if a Slammer CU warmstarts, the recovery may fail. Multiple Slammer CUs may go offline due to repeated software faults, and may not be able to return online.
		[14317229] Added the ability to accept new Bricks and assign them to a Storage Domain to the Monitor - Hardware - Bricks screen.
		[14312940] The FLT, ST, and ACT LEDs on a Slammer CU may show incorrect status after a software update, system start up, or Slammer CU maintenance. The fix updates the LEDs to the proper status if the Slammer details are selected.
		[14303555] Resolves an issue for the basic web interface that may cause some scanners to report a positive for CVE-2011-3192 Request-Range vulnerability.
		[14276163] On large Axiom systems with 3 or 4 Slammers, the processing of Slammer Statistics may cause repeating periodic Pilot software faults.
		[14269001] Certain advanced Brick maintenance operations using the console function may cause the RAID controller to experience a software fault due to the load causing heartbeat timeouts.
		[14204680] Under some conditions, the request to generate a firmware debug dump of the iSCSI interface may fail due to duplicate debug requests.
		[14177287] The average throughput statistics counters for Read and Write in the Slammer Protocols statistics are reversed, showing the "read" traffic for Writes and the "write" traffic for Reads.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
		[14172759] If a Slammer CU replacement motherboard requires a PROM update, a failure in parsing Power On Self-Test results on the new motherboard may result in the PROM not being updated.
		[14154953] Pilot CU replacement with the A1811 Enhanced Pilot may fail because the internal network has not stabilized before the utility that checks and updates the new Pilot CU begins to run.
		[14138379] Slammer CUs may warmstart as Network Time Protocol information is being updated.
		[14128575] Email alerts may fail if the list of addresses includes any that are not in the organization email domain.
		[14126312] A grace period of 60 seconds has been added to the Pilot management link down detection. If the external management link status is restored within 60 seconds, the Pilot CU will not fail over to the standby pilot.
		[14121451] Adds Administrator Alert for link down conditions on the Private Management PMI network ports on the Pilot CUs.
		[14086372] Adds a new field, AllocatedCapacityInBytes, to the output of Axiom CLI information for LUNs and Clone LUNs.
		[14086151] Long LUN names are truncated to approximately 26-28 characters in generated reports.
		[14039336] Under certain conditions, Fibre Channel disk write errors and the corresponding normal internal sector repair operations performed could cause data inconsistency.
		[14008432] If the Slammer CUs do not shut down cleanly on the disruptive upgrade from Release 4.x, they may experience repeated software faults during the system startup on the R5 release.
		[13815063] SNMP traps for Resync Required and Bandwidth shaping events from the MaxRep Engine do not generate event driven logs from the Axiom system.
		[13766896] As Storage Enclosures (Bricks) are added to existing systems, the Axiom Serial Number may not be updated correctly in the Vital Product Data stored in the Enclosure Midplane.
		[13765570] After a Slammer CU Failover/Failback sequence, LUNs may go offline with the Lost Data flag set due to a timing issue in recovering the metadata for the LUNs.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
5.3	All	[13765523] If the Active Pilot CU cannot reach the configured gateway on the external management interface, but the Standby Pilot CU is able to, the Active Pilot will fail over to the Standby Pilot and generate an administrator alert.
		[13765310] If multiple hosts access the same Axiom system, and there is a improperly duplicated initiator identity in those hosts, LUNs mapped to one host may become mapped to other hosts with the same initiator identity.
		[13761058] The help menu option has been removed from system alert screens.
		[13757918] The ability to select and delete multiple LUNs at the same time using the GUI has been added.
		[13747868] Replacing the network interface module (NIM) in each Slammer control unit (CU) within five minutes of each other can bring down the host connections.
5.2	All	[13757970] In the Pillar Axiom Storage Services Manager, column sorting the components in the Brick hardware details view does not sort the drives in numbered order.
		[13762034] In the Pillar Axiom Storage Services Manager, the Tools menu contains no items, even though it appears to be active.
		[13759900; 13761702] Sometimes, the Pillar Axiom Storage Services Manager GUI consumes too much memory and CPU cycles after being left running for several days.
5.0	SAN	[13762643] During LUN creation, the Storage Class field does not default to a specific Storage Class in the Quality of Service (QoS) tab.
		[13762633; 13759853] When upgrading from release 4.x to release 5.0, even though the Quality of Service (QoS) settings for the LUNs are correct, all LUNs show "Custom" as the performance profile.
		[13762691] Due to an incompatibility issue between the firmware on legacy Fibre Channel (FC) Bricks in releases below 4.3.0 and release 5.0, all system upgrades to release 5.0 need to first be upgraded to release 4.3.15 or higher as an intermediate step. If this intermediate step is not performed, any legacy FC Bricks on the system end up Critical with the possibility for data loss.
		[13762042; 13761194] In certain circumstances, a Fibre Channel Brick can get a stuck command on one RAID controller when a Pre-Emptive Copy (sometimes called <i>copy-away</i>) completes. This situation can result in I/O to that controller being halted, which can result in logical volumes going offline.

Fixed in Release	Slammer type	Pillar Axiom server issue or impact
4.3	SAN	[13762274] When mapping a LUN to Host Groups, the GUI lets the user attempt to change the LUN number mapping for a given host in the Host Group, although this is not allowed. All hosts in a given Host Group receive the same LUN number mapping.
		[13759258] A Pillar Axiom system might experience a software fault if configured to enable iSNS, but with an invalid iSNS server IP address.
	All	[13753680] During a warmstart, attempts to reinitialize the internal Ethernet controller can sometimes fail, which results in a Slammer control unit (CU) failover.
		[13753622] Unable to parse and see Brick S.M.A.R.T. data using the Pillar Axiom Statistics Tools.
	SAN	[13756313] When isolating a replication pair using the command <code>axmrepsan isolate_replication_pair <i>pair_name</i></code> , sometimes the replication pair can enter the AWRY state, which is verified in the output generated by the command <code>axmrepsan get_replication_pair_status</code> .

9 Additional Notes

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

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

9.1 Host Queue Depth on SAN Hosts

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

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

9.2 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.

9.3 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.

9.4 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.

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

[13747833] During the Guided Maintenance (GM) of Slammer CU components, the system attempts to flush all pending I/O writes to storage so that the system can go into write-through mode. Going into write-through mode ensures that no data is in cache should a catastrophic error occur. If there is heavy write activity during this time, an alert may be generated in the GUI warning the administrator that the system cannot properly prepare for the Guided Maintenance.

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

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

[13746134] Under certain circumstances, the Microsoft iSNS Server v3.0 may add or remove members of the Pillar Axiom's discovery domain without notifying the Pillar Axiom system. If iSNS access control is enabled, the missing notifications can cause the Pillar Axiom iSCSI target to accept or reject iSCSI initiator logins when it should not.

To prevent this problem from occurring, follow these guidelines:

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

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

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

9.7 HP-UX HBA Connections to Pillar Axiom Systems

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

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

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

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

9.8 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.

9.9 LUN Assignment When the System Restarts

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

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

9.10 Blacklisting Local Drives on RHEL4 Platforms

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

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

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

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

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

9.11 iSCSI Software Initiator May Have Two Names Associated With It

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

9.12 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 Oracle Customer 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.

9.13 Uploading Software Update Packages over Slow Connections

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

9.14 When Updating the Software

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

9.15 Software Update Process Waits for all Tasks to Complete

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

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

- Schedule updates when it is known that long-running tasks will not be running.
- Before starting a software update, be sure all tasks are complete by clicking **Tasks** on the lower right corner in the GUI.
- During the software update, review the tasks that are running. However, before you cancel those tasks, contact Oracle Customer Support for assistance in determining which, if any, tasks can be canceled.

9.16 Non-Disruptive Software Updates

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

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

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

9.17 WWN Designation Changed in Release 2.0

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

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

[13759444] Starting with release 4.3, the Fibre Channel (FC) World Wide Name (WWN) for the Series 3 Slammer that contains a SAN 8 Gb/s FC network interface module (NIM) is derived from the 8 Gb/s FC HBA itself. The WWNs for both Slammer CU ports are printed on a label located on the faceplate of the HBA. They are also displayed in the GUI.

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

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

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

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

9.19 Dynamic LUN Expansion Not Well Supported on Linux

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

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

9.20 Status of Logical Volumes

Beginning with release 5.0, the GUI displays the overall status of the Pillar Axiom system and of its hardware as adjacent icons on the far left side of the status bar. However, these icons do not reflect the status of the logical volumes and their data.

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

The system alert icon (flashing yellow triangle) is adjacent to the hardware status icon and indicates whether any system alerts exist, including those related to logical volumes. When this indicator is present, navigate to **Monitor > System Alerts** to check the details of the alerts.

To check the status of the logical volumes, navigate to **Configure > Storage > San > LUNs**.

9.21 Changing the Time on a Pillar Axiom System

Tip! Because the Pillar Axiom policy controller is Linux-based, Pillar recommends that you not use Windows servers as NTP servers. For more information, refer to the following Microsoft TechNet article and, if desired, contact Oracle Customer Support.

<http://technet.microsoft.com/en-us/library/bb463171.aspx>

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

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

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

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

9.22 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 creates the following system alert:

```
SnapLUNStorageFillingButCannotGrow  
("Extra space for Clone LUNs has reached maximum and is nearly full")
```

When repository usage crosses the 99% threshold, the system creates this same alert. When the repository is 100% consumed, all clones are invalid. 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, 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.

9.23 Deleting Clones From the Youngest to the Oldest

[51125] When several clones are deleted at the same time using the GUI, the process can take about seven minutes for each clone. The length of time it takes to delete a clone is related to the time it was created and how it relates to the clone storage space being used. The best approach to deleting clones would be to choose the youngest clone first and work your way back to deleting the oldest.

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

[13762598] Deleting the youngest clone may take a long time because this action requires reassignment of data to ensure that the older clones (if any exist) will have consistent data.

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

9.24 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.

9.25 Reassigning a Logical Volume to Another Slammer or Control Unit

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

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

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

9.26 System Scaling

You can increase storage capacity on a system by adding components and increasing assigned capacities.

As of release 5.0, you can decrease storage capacity of a system through the use of Storage Domains. Though you cannot decrease the capacity of the primary Storage Domain, you can decrease the capacity of a non-primary domain by removing all data from that domain and then use the **Monitor > Hardware > Bricks > Remove Brick** action in the GUI to remove the Brick from the system.

9.27 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.

9.28 Cause of Degraded Status for a Volume Has Changed

[13751747] Beginning with release 4.0, a failed, missing, or pulled drive will not by itself cause a logical volume to become degraded. Only an inability to write to the mirror in a doubly redundant volume will result in a degraded status.

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

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

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

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

9.30 Replacing Hardware Components

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

9.31 Replacing or Upgrading the Slammer Requires Cloning the Slammer

[13766221] The non-volatile random access memory (NVRAM) in a Slammer storage controller contains the world wide node name (WWNN), which uniquely identifies the Slammer chassis. To upgrade or replace an existing Slammer, this WWNN must be duplicated onto the replacement Slammer. This process of checking, modifying, and updating the Slammer WWNN is called *cloning*. Contact Oracle Customer Support to clone a Slammer or to check the cloning on a Slammer.

9.32 Alternative to Guided Maintenance Identify Step for Pilots

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

You can identify a Pilot CU by matching the serial numbers reported in the GUI with the labels on the Pilot CUs. However, if the Pillar Axiom system is local, the easiest method is to use a combination of pinging the Pilot CU and then disconnecting the management cable to that CU and re-pinging the disconnected CU to confirm its identity.

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

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

9.34 Reuse of Pilot Control Units Can Cause Problems

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

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

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

9.35 Differentiate Between FC RAID Bricks and FC Expansion Bricks.

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

9.36 Replacing a Drive

When replacing a drive, always use a new one from Oracle.

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

Important! You should contact Oracle Customer Support before pulling a drive for test purposes.

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

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

Contact Oracle Customer Support for a new replacement drive.

9.37 Moving Drives

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

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

9.38 Reseat Drives before Powering On New or Replacement Bricks

Important! As a precautionary step, do not power on a new or a replacement Brick until you have completed the following steps.

After adding the Brick to the system, one or more drive latches might not be in full contact with the Brick chassis midplane. A drive latch may appear to be fully latched, but sometimes the drive is not making good contact with the midplane. With poor contact, the drive will fault and its status will be displayed as Critical (because it is missing).

To prevent loose drives, visually inspect each drive to verify that they are fully seated. If a drive is not fully seated, either or both of the following will be true:

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

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

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

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

9.39 Testing Multiple Drive Failures

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

9.40 ACT LEDs on Drives Can Blink When Inactive

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

9.41 Replacement of Brick Storage Enclosures

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

9.42 Adding a Brick Generates Error and Warning Messages

When you add a Brick storage enclosure to a Pillar Axiom system, before powering on the Brick, carefully check the cabling to ensure the cabling is correct.

When you power on the Brick, the system begins the process of bringing the Brick online. While the system is bringing the Brick online, you will see multiple instances of the message:

Topology Discovery Task.

Note: If the `Topology Discovery Task` does not complete after approximately 10 minutes, contact Oracle Customer Support.

After you power on the Brick, the system creates a system alert. This alert allows you to assign the Brick to a particular Storage Domain.

After you assign the Brick to a Storage Domain, check the status of any existing LUNs. If any LUN goes Offline or Partial Offline, or if a system alert indicating Pinned Data is generated, power off the new Brick and contact the Support Center for assistance.

While the system is bringing the Brick online, you might also see a series of error and warning messages similar to the following:

- 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 goes from red to yellow to green. After the system completes the process, the Brick shows a Normal status and removes all alerts related to adding the Brick. If any alerts remain, contact Oracle Customer Support.

Tip! After the Brick has been successfully added, verify that the system total capacity has been updated to include the storage provided by the new Brick.

9.43 Testing RAID Rebuild and Simulated Drive Fault

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

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

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

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

9.44 Brick Issues Can Cause a Slammer to Warmstart

[13747255] The Array Manager software component in Pillar Axiom Slammers requires a quorum of successful I/O for its metadata processing. Sometimes a quorum cannot be met because some Bricks return a Busy state due to underlying issues on the Bricks. After a number of unsuccessful retries, the Array Manager can fail a health check and cause the Slammer to warmstart.

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

9.45 Testing Brick Power Loss

Oracle ACS can assist in testing power loss to Brick storage enclosures.

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

[13746893] In any given Brick storage enclosure, if there are enough drive failures that cause a data LUN on that storage enclosure to fault, care should be used in recovery.

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

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

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

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

9.47 Creating Logical Volumes Immediately After Replacing a Failed Drive

[13751476] When a failed drive in a SATA Brick storage enclosure is replaced, the system copies the data temporarily stored on the spare drive to the drive replacement. This write operation is called *copyback*. When the copyback operation starts, the unused storage on the affected Brick temporarily becomes unavailable for new allocation. This condition is necessary to guarantee that the unused space is correctly reconditioned.

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

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

9.48 System Configuration Vulnerability in Single-SATA Brick Systems

[13752315] When a factory-fresh, single-SATA Brick system is first powered up, the system configuration database (which is the Pillar Axiom data store containing many system settings) will be configured on that Brick. This database will be doubly-redundant with both instances residing on the same Brick (on two data LUNs). If additional Bricks are added later, the system configuration database will not be migrated to locate the different copies on separate Bricks for higher protection against single-Brick failure.

Important! The system configuration database is still vulnerable to that original, single SATA Brick going offline, as will be user data. [13765974] As of release 5.0, if this database goes offline, the system will not perform an emergency restart or shutdown, but you will not be able to modify or access the information for any resource in the system configuration database. Any request that requires access or modification will fail. To recover, contact Oracle Customer Support.

To mitigate this vulnerability, as of release 5.0, you can use Storage Domains to cause the system to migrate the redundant instances of the system configuration database onto separate Bricks by performing the following general actions:

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

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

Tip: For assistance in determining which Bricks contain the system configuration database, contact Oracle Customer Support.

9.49 Replacing a Slammer Motherboard Can Cause Several Status Changes

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

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

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

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

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

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

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

Important! Do not run Slammer diagnostics without explicit instructions and assistance from Oracle Customer Support.

9.50 Testing Failure Recovery from Loss of Slammer Power

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

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

9.51 Uploading Empty Files Through the GUI Is Not Allowed

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

9.52 DHCP Behavior on the Pilot

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

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

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

DNS needs to be enabled for the Pilot management console to be able to make an outbound connection, such as for Call Home services.

Note: Your ability to log in to the Pillar Axiom system using the system name rather than its IP address is not dependent on the DNS setting in the Pilot.

9.53 Changing Slammer Port IP Addresses from Static to Dynamic

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

9.54 VSS Provider Event Numbers Incorrectly Mapped to Descriptions

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

9.55 Disabled/Excluded Slammer States

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

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

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

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

Note: Contact Oracle Customer Support for assistance in recovery of Slammer CUs that have been marked Disabled.

9.56 Slammer Warmstart and Startup Failure Handling

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

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

Contact Oracle Customer Support for recovery assistance for any Slammer CU that is Excluded or Disabled.

9.57 Slammer Additions and Replacements

When adding a new Slammer (Series 4) to a Pillar Axiom system or replacing an existing Slammer, be sure that the Axiom system is running release 5.3 or higher.

9.58 Hardware Lockout Due To Repeated Power Cycling

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

The current thresholds for repeated power cycles are:

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

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

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

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

9.59 Powering Off a Pillar Axiom System

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

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

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

9.60 Battery Removal

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

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

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

9.61 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 Oracle Customer Support for assistance in checking the state of the batteries or for a replacement.

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

9.62 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 Oracle Customer Support.

9.63 Hardware Component States

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

9.64 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.

9.65 OpenSSH Has Some Vulnerabilities

Pillar addresses security vulnerability in these ways:

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

- mime_module
 - mpm_worker_module
 - rewrite_module
 - setenvif_module
 - so_module
 - ssl_module
- Pillar Axiom systems will support the latest version of OpenSSH and SSL in an upcoming release.

9.66 Shutting Down a Pillar Axiom System

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

9.67 Login to Oracle EM Plug-In Fails

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

9.68 SAN Host Associations and LUN Mappings

9.68.1 Definitions

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

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

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

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

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

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

9.68.2 Managing Unassociated and Associated Hosts

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

9.68.2.1 Moving an Initiator From an Unassociated Host

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

9.68.2.2 Moving an Initiator From an Associated Host

When the administrator moves an initiator from one associated host to another, the following actions occur:

- The mappings of the previous associated host are removed from the initiator.
- The initiator is mapped with the mappings of the new associated host.

The previous associated host (even if it has no initiators left) is left with its LUN mappings in place. No mappings associated with the previous associated host are moved to the new associated host. The previous host can be manually deleted, if it's no longer needed.

9.68.2.3 Deleting a Host

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

9.68.2.4 Renaming a Host

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

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

9.68.3 Managing APM Hosts

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

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

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

9.68.3.1 Moving an Initiator From an Unassociated or Associated Host

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

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

9.68.3.2 Moving an Initiator From Another APM Host

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

9.68.3.3 Combinations

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

9.68.3.4 Removing an Initiator From an APM Host

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

9.68.3.5 Deleting an APM Host

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

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

9.68.3.6 Renaming an APM Host

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

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

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

9.69 Host Mappings Can Disappear Under Certain Circumstances

[13754097] If a SAN host with associated initiators was defined in the Pillar Axiom system before release 2.5 and that host subsequently has initiators added and removed in a single operation (either using the GUI or the PerformAssociateInitiatorsToHost CLI request), the defined SAN host and the associated initiators may disappear.

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

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

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

9.70 Overloaded Systems Can Result in Multiple System Warmstarts

[13748319] The Pillar Axiom Pilot software is multi-threaded, and the number of threads available for handling system events and user-requested operations is limited. When a large amount of system activity requires Pilot intervention, this maximum number of threads can become fully utilized, leaving no threads available for user-requested operations.

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

9.71 IPStor Fails To Recognize LUNs After a Full System Upgrade

[13749164] FalconStor IPStor systems do not recognize Pillar Axiom LUNs after a system upgrade from a Pillar Axiom 500 to a Pillar Axiom 600. IPStor systems recognize only the Pillar Axiom 500 as a valid target. To resolve this situation, request from your IPStor vendor the code update necessary for compatibility with the full Pillar Axiom product line.

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

9.72 Getting Accurate SNMP Filesystem Performance Metrics

[13753980] If an SNMP request is made within 60 seconds of the last request of the same type, the system returns the same information as it did in the prior request. To ensure you get the latest information, wait more than 60 seconds between identical SNMP requests.

10 Technical Documentation Errata

The following entries describe information in the technical documentation that could not be corrected in time for the current release.

10.1 Pillar Axiom Administrator's Guide

- At the top of page 27, replace Step 3 with the following:
 - 3 Click the **Management Software** tab.

- On page 58, replace Steps 3 and 4 with the following:
 - 3 Click the **Utilities Software** tab.
 - 4 Click the **Download Axiom SNMP MIP** link.

- On page 86, delete the Note that begins with "Solid-state drive (SSD Bricks"

- On page 147, replace Steps 1 and 2 with the following:
 - 1 Start a web browser from your workstation.
 - 2 Specify the IP address of the Pilot management controller or the name of the Pillar Axiom system as the address to open.
 - 3 Click the **Utilities Software** tab.
 - 4 Under **Axiom Plug-Ins**, click the link to download the VSS Provider for your operating system type.

Valid VSS providers:

 - **Download Pillar Axiom VSS Provider - 32 bit**
 - **Download Pillar Axiom VSS Provider - 64 bit**

- On page 164, replace Steps 1 and 2 with the following:
 - 1 Start a web browser from your workstation.
 - 2 Specify the IP address of the Pilot management controller or the name of the Pillar Axiom system as the address to open.
 - 3 Click the **Utilities Software** tab.
 - 4 Under **Axiom Plug-Ins**, click the link to download the VDS Provider for your operating system type.

Valid VDS providers:

 - **Download Pillar Axiom VDS Provider - 32 bit**
 - **Download Pillar Axiom VDS Provider - 64 bit**

- For the software and firmware update content on pages 167 through 175, refer to Section 3.5 for the replacement text.
- On page 232, replace the third paragraph with the following:

While the system is in a shutdown state, the only actions you can perform are to display system status, to restart the system, and to power down the hardware.
- On page 232, move the **Important!** Notice to the end of the procedure.
- On page 269, replace the third paragraph with the following:

Printing a trending chart enables you to record the trends displayed in the chart for later use.
- On pages 325 and 327, delete the entry for **Number of Paths**.
- On page 552, delete the entry for **Collection Reason**.
- On page 563, delete the entries for **Target** and **Reason**.
- On page 574, add the following item to the Software Update Options:

Ignore system alerts
Select this option to update the software and ignore any System Alerts.
- The **Support > Utilities** page has been removed from the Pillar Axiom Storage Services Manager graphical user interface (GUI). The software that was available from the Utilities page has been moved to the Pilot management controller web interface.

Because of that change in the GUI, delete the following topics:

 - Page 578, Utilities Overview Page
 - Page 580, Virtual Disk Service (VDS) Page
 - Page 581, Volume Shadow Copy Service (VSS) Page

10.2 Pillar Axiom CLI Reference Guide

- [14667547] For the `-modify` syntax of the following commands, change the `-name` option to `-newname`:

```
account
brick
event_notification
host_group
job
san_host
snmp_host
system
ups
```

- On page 220, under the section for `-add`, in the first paragraph, replace the second sentence (which starts with “This command is used …”) with the following text:

To create a SAN host entry for a SAN host that is not running Pillar Axiom Path Manager (APM), the `-add` subcommand uses the World Wide Names (WWNs) of HBA FC or iSCSI ports for the host.

WARNING: *Do not use the `-add` subcommand for hosts that are running APM. When a host that is running APM is connected to the Pilot, APM automatically creates an entry in the system for that host. Using the `-add` subcommand for hosts that are running APM might corrupt the internal system tables and result in an error.*

- On page 288, replace the `-logBundle` option for both the `-download` and `-send` subcommands with the following:

```
[-logBundle log_bundle_id_or_fqn [,log-bundle-id-or-fqn]...]
```

- On page 291, in the description of the `-download` subcommand, replace the first three paragraphs with the following:

Downloads one or more tar files with statistics from the Pillar Axiom system to your administrative workstation. Transmit the files as requested by Oracle Customer Support.

Valid options:

```
-logBundle
```

Identifies one or more log bundles by either the ID or FQN. If you do not specify a log bundle, the most recent log bundle is downloaded.

- At the bottom of page 291, under the `-file` option of the `-download` subcommand, add the following after the existing sentence:

If multiple bundles are downloaded, each will have an integer suffix appended to the file name.

- At the top of page 293, replace the first sentence, which begins with “Enter the ID or FQN ...,” with the following:

Enter the ID or FQN of one or more specific Call_Home bundles to send.

10.3 Pillar Axiom 600 Hardware Installation Guide

- On page 41, Table 8, for the second, third, and fourth rows, replace the values in the third column (“Maximum number of FC Bricks”) with 64.

(In other words, all Pillar Axiom 600 systems that have more than one Slammer can now support up to 64 FC Bricks. The number of SATA and SSD Bricks remain the same, as listed.)

- On page 41, replace the entire paragraph that starts with “Pillar Axiom 600 systems support up to 32 FC Bricks ...” with the following sentence:

Pillar Axiom 600 systems support up to 64 FC Bricks, or 64 SATA Bricks, or 32 SSD Bricks or any combination of these three types, provided there are no more than 64 total Bricks, 64 FC Bricks, or 32 SSD Bricks in the system.

- On page 45, replace the sentence that starts with “Pillar Axiom 600 systems with two or three Slammers ...” with the following:

Pillar Axiom 600 systems with two or three Slammers support up to 64 FC Bricks.

- [13766986] On page 97, Table 14, replace the content of the first cell in the fourth row with the following:

All Pillar Axiom systems that have more than three Slammers or more than 32 Bricks require optical private interconnect (PI) cabling for higher performance, better stability, and easier fault location.

If you have an older Pillar Axiom system that supports only copper cabling and has version 1 private interconnect modules (PIMs), consult Oracle Customer Support to upgrade the Pillar Axiom system to support more than 32 Bricks and two Slammers. However, no upgrades are provided to convert version 1 copper based cabling to version 2 optical cabling. If Bricks are added later, a re-cabling may be necessary to ensure that there are no more than eight Bricks on any Brick string.

- On page 108, replace the second bullet in the Prerequisites list with the following:

The external switch ports must be configured for auto-negotiation, if the Pillar Axiom 600 system is also configured for auto-negotiation. If the external switch ports are manually set to a certain speed or a duplex setting, the Pilot must be set to match those settings.

- On page 116, replace the paragraph that begins with “The Pulizzi PC3365 shown in the following figure requires ...” with the following paragraph:

The Pulizzi PC3365 shown in the following figure requires 230/400 VAC WYE three phase input power. The output at each outlet is always 230V. This PDU is only available in the United States.

10.4 Pillar Axiom 600 SSF Cabling Reference

On page 49, in the eighth bullet from the bottom, which begins with “FS-4 and FS-5 are reserved ...”, change the expression “3- and 5-Slammer configurations” to “3- and 4-Slammer configurations”.

[13766986] On page 101, Table 40, replace the content of the first cell in the fifth row with the following:

All Pillar Axiom systems that have more than three Slammers or more than 32 Bricks require optical private interconnect (PI) cabling for higher performance, better stability, and easier fault location.

If you have an older Pillar Axiom system that supports only copper cabling and has version 1 private interconnect modules (PIMs), consult Oracle Customer Support to upgrade the system to support more than 32 Bricks and two Slammers. However, no upgrades are provided to convert version 1 copper based cabling to version 2 optical cabling. If Bricks are added later, a re-cabling might be necessary to ensure that no more than eight Bricks exist on any Brick string.