Pillar Axiom

System Architecture Overview

ORACLE

PILLAR AXIOM

Part Number: 4420-00029-1300
Pillar Axiom release 5.0
2011 October

Table of Contents

List of Figures

List of Tables

Preface

Audience

This documentation is intended for individuals who plan, purchase, and implement enterprise storage systems.

Related Documentation

The following Pillar Axiom technical documentation may help you succeed in the use of this document.

- *Pillar Axiom Administrator's Guide*
- *Pillar Axiom Statistics Tools User Guide*
- *Pillar Axiom SnapDelta FS Reference Guide*
- *Pillar Axiom CIFS and NFS Multi-Protocol Planning Guide*
- *Pillar Axiom iSCSI Integration Guide for Windows Platforms*
- *Pillar Axiom NDMP Integration Guide*
- *Pillar Axiom Service Guide*
- *Pillar Axiom SMIProvider Reference*
- *Pillar Axiom SSF Cabling Reference*
- *Pillar Axiom Support and Interoperability Guide*
- *Pillar Axiom Glossary*

Access Documentation

Technical documentation (including installation, service, cabling, integration, and administration guides) for Oracle's Pillar Axiom 600 storage system is available from several sources.

Pillar Axiom Storage Services Manager	Log in to your Pillar Axiom system. Navigate to the **Support** area in the Pillar Axiom Storage Services Manager and select the **Documentation** link.
Pillar Axiom HTTP access	For Pillar Axiom systems running release 5.0 (and higher) software, point your browser to `http://system-name-IP/documentation.php`, where *system-name-IP* is the name or the public IP address of your system.
Internet	Customer support portal (http://support-portal.pillardata.com/csportal/login.seam). Log in and click **Documents** in the left navigation pane.
Product CD-ROM	Insert the Technical Documentation CD-ROM (came with your Pillar Axiom system) into the CD player and open the `DocMenu` PDF. **Tip:** To search all technical documents on the CD-ROM, click **Search all PDFs** in the top right corner.

Pillar Contacts

Table 1 Contacts at Pillar Data Systems

For help with...	Contact...
Error messages, usage questions, and other support issues	US and Canada: 877-4PILLAR (1-877-474-5527) Europe: +800 PILLAR FS (+800 74 55 27 37) Asia Pacific: +1-408-518-4515 South Africa: +0 800 980 400 Have your system serial number ready. support@pillardata.com Pillar Customer Support (http://support-portal.pillardata.com/csportal/login.seam)
Training (custom or packaged)	Training and Education (http://www.pillardata.com/support-education/training/)
Professional services and inquiries	globalsolutions@pillardata.com

Table 1 Contacts at Pillar Data Systems (continued)

For help with...	Contact...
	Global Solutions (http://www.pillardata.com/support/professional-services/)
Sales and general contact information	Company contacts (http://www.pillardata.com/company/contact)
Documentation improvements and resources	docs@pillardata.com Technical documents (http://www.pillardata.com/techdocs) (Log in with your username and password, and select Documents.)

CHAPTER 1

Welcome to the Pillar Axiom System Architecture

Foreword

This version of the *System Architecture Overview* describes Pillar Axiom system release 5.

This document describes both network attached storage (NAS) and storage area network (SAN) features, but only the SAN features are supported in release 5 of the Pillar Axiom software.

Product Overview

The Pillar Axiom system is a complete and integrated full-featured network storage system.

Pillar Axiom systems offer a unified SAN and NAS platform that combines:

- Innovative performance characteristics
- The unprecedented ability to enable administrators to assign storage resources (capacity, CPU, and cache) based on business requirements or conditions
- Consistent Quality of Service (QoS) tools to manage application I/O contention
- Serviceability
- Data migration
- Distributed RAID technology

And many other industry-leading features.

Feature Overview

Features of the Pillar Axiom storage solution include:

Policy-based Quality of Service (QoS)	The Pillar Axiom system employs intelligent storage device management to provide QoS capabilities that yield a very high utilization of the storage resources in the entire system. This QoS functionality enables the system administrator to prioritize data according to the importance of the applications and the required performance for each data store.
Storage Domains	Storage Domains provide a subset of a virtual storage pool comprised of a grouping of physical Bricks. Storage Domains are typically used to provide specific allocation or security features for a collection of logical volumes (filesystems or LUNs).
Storage Classes	Storage Classes allow you to specify the preferred storage media to use for a logical volume (filesystem or LUN). This makes it possible to configure QoS for each drive type that the Pillar Axiom system supports.
Optimized capacity management	With its built-in capacity planning tools, the Pillar Axiom system achieves the most efficient use of available storage capacity.
Advanced backup and recovery methods	A variety of data protection tools are available in the Pillar Axiom system to provide backup and restore capabilities to meet a wide range of recovery point objective (RPO) and recovery time objective (RTO) requirements.
Replication	Replication of both NAS filesystems and SAN volumes facilitates automated backup and restore of mission critical data to ensure disaster recovery and business continuance.
Storage Profiles	Numerous storage profiles are available to define Quality of Service (QoS) attributes such as performance level and RAID characteristics for individual logical volumes (filesystems or LUNs). Storage profiles help storage administrators maximize performance.
Scalability	Pillar Axiom systems are designed to scale as business needs expand. Built-in hardware and

	software flexibility make it easy to add capacity, cache, or CPU power.
NAS and SAN support	NAS and SAN systems can coexist and share the same storage pool.
Flexible storage reallocation	Storage can be allocated or reallocated to different Storage Classes or Storage Domains to meet performance and business needs without disrupting normal operations.
Alerts and event management	Easy-to-use software facilitates monitoring and troubleshooting the Pillar Axiom system.
System monitoring	An intuitive graphical user interface (GUI) provides complete information about system performance at a glance.
Statistical reports	A sophisticated reporting interface facilitates gathering and analyzing all types of system data that administrators can use to optimize system performance.
Interoperability	Pillar Axiom systems can be used with a wide variety of popular hardware and software.
Storage volume expansion	Filesystems and LUNs can be expanded to meet growing business needs.

CHAPTER 2

Pillar Axiom Hardware Overview

Pillar Axiom System Components

The Pillar Axiom system provides SAN and NAS connectivity to a common pool of storage.

The system is modular. The three major system components are:

- Pilot management controller
- Slammer storage controllers
- Brick storage enclosures

Slammers and Bricks communicate using a highly redundant Storage System Fabric (SSF).

Slammers and Bricks have:

- Redundant power supplies and fans
- Front and back LEDs that provide system and component identification and status
- Built-in RAS: reliability, availability, and serviceability
- Field replaceable units (FRUs)

The system components fit into a Pillar-supplied rack or a standard D-class 4-post 19-inch rack. Larger configurations require multiple adjacent racks.

The Pillar Axiom system continually monitors all hardware components for proper operation and fault status.

The following figure shows the interactions between the Pillar Axiom system components.

Figure 1 Interactions among Pillar Axiom components

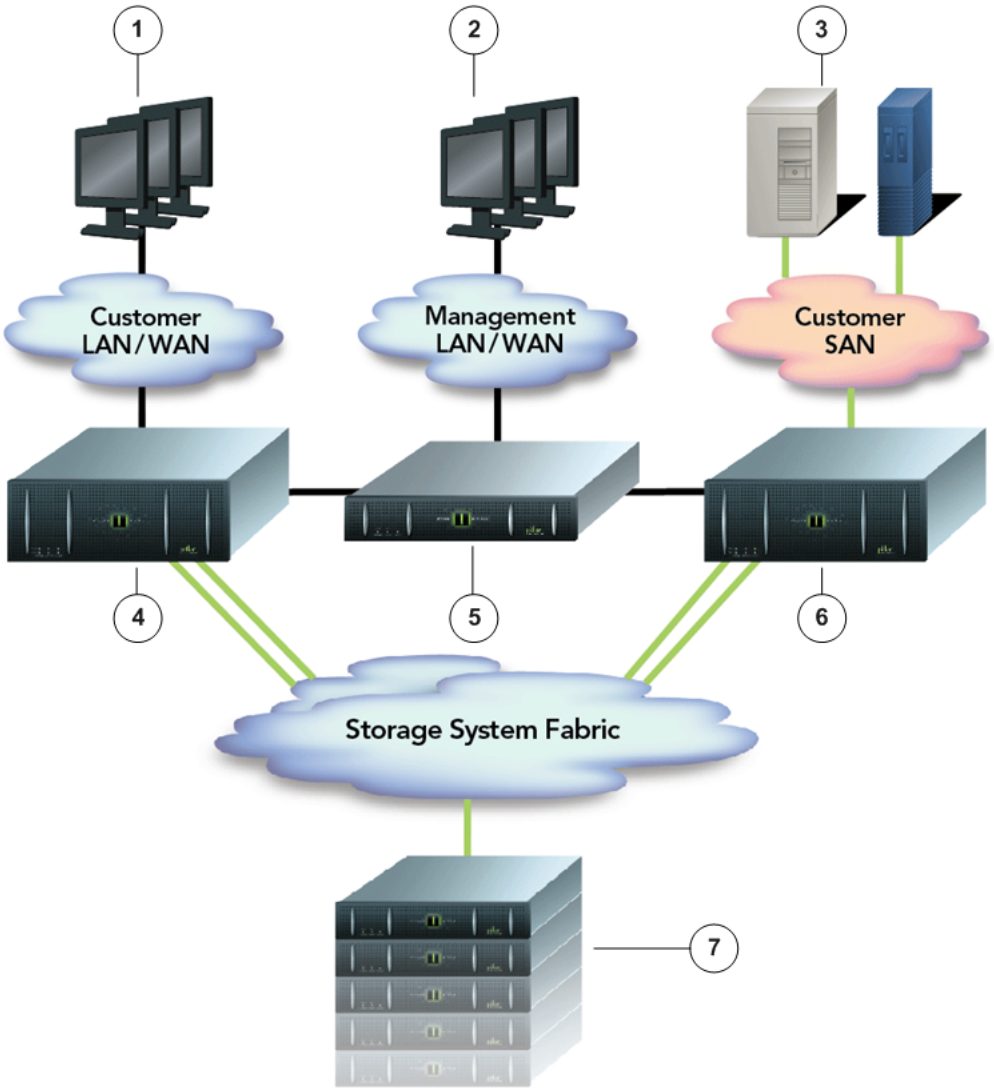

Legend

1 NAS clients (NFS or CIFS)	5 Pilot management controller
2 Management clients	6 SAN Slammer storage controllers
3 SAN clients (FC or iSCSI)	7 Brick storage enclosures
4 NAS Slammer storage controllers	

Storage System Fabric (SSF)

The Storage System Fabric (SSF) consists of a switched Fabric with multiple redundancy that carries all data traffic among Slammer storage controllers and Brick storage enclosures.

The SSF is defined as:

The protected Fibre Channel fabric internal to Pillar Axiom systems that interconnects Bricks and Slammers. The SSF enables communication within the Pillar Axiom system so that all Slammers can connect to any of the Bricks. The SSF provides redundant paths for increasing reliability.

Each control unit (CU) within a Slammer has a private interface module (PIM) that contains a Fibre Channel switch. The Fibre Channel ports on the switch connect the CU to the SSF. The SSF utilizes either a 2 Gb/s copper or a 4 Gb/s optical Fibre Channel interface to provide sufficient bandwidth.

Every port of the switch can simultaneously service data I/O operations. The Fabric Manager controls all data transfers throughout the SSF.

Pilot Management Controllers

The Pilot is an out-of-band management system that directs and manages all system activity.

Each Pillar Axiom system has one Pilot, which contains the following hardware:

- Redundant Ethernet interfaces connected to the customer network for administration and support of the Pillar Axiom system.
- Redundant Ethernet interfaces connected to the Pillar Axiom private internal network for communication with the Slammers and for coordination between the two Pilot control units (CUs).
- Serial ports connecting the two Pilot CUs as a backup for managing the active or standby status.
- Storage for software, logs, and configuration cache.

Figure 2 Pilot components

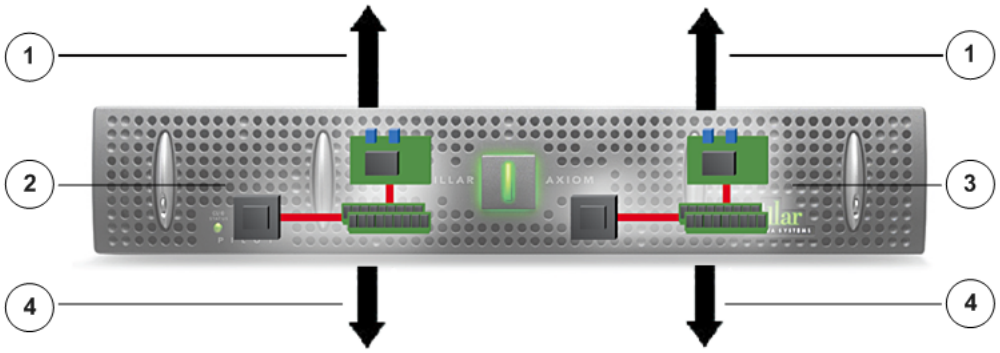

Legend	
	1 To external management network
	2 Control unit CU0
	3 Control unit CU1
	4 To Slammers

The two independent Pilot CUs operate in active or standby mode. The standby Pilot continually monitors the health of the active Pilot and will promote itself to active status if necessary.

The Pilot is not connected to any data path interface and has no direct access to user data.

The Pilot provides access to all management user interface functions for the Pillar Axiom system:

- Graphical user interface
- Command line interface
- SMI-S Provider interface
- Oracle ASM interface
- Microsoft VSS Provider interface

The Pilot provides the configuration management interface to Pillar Axiom Path Manager (APM) clients and replication pairs.

The Pilot provides alerts, log collection, and automatic log transmission.

The Pilot provides management for booting, replacing, updating, or recovering Slammer or Brick resources.

Pilot Functional Description

The Pilot configures and manages storage through Quality of Service (QoS) policies that allocate hardware assets based on need and service level.

The Pilot also:

- Manages backups.
- Manages restoration of backed up data sets.
- Supplies system monitoring, notification, and reporting services.
- Monitors SAN host clients.
- Manages replication.

The Pilot directs the operation of the Slammer based on user settings from the GUI or CLI. A private management interface (PMI) is shared between the Pilot CUs and the Slammer CUs so they can access information from the Slammers.

A heartbeat running between all of the Pilot and Slammer CUs monitors the status of these CUs. In case one CU fails, the other CU becomes active and takes over management. The heartbeat utilizes both the private Ethernet and a backup serial cable.

The Pilot also monitors any replacement Pilot CUs that are added to the system and compares them to the active CU. It revises the software on the new Slammer, Pilot, or Brick CU so that it matches that of the current, active CU.

Pilot Software Components

The Pilot management controller includes the following software components:

Data Mover API — The Data Mover Remote API provides the interface to the data-mover engine located on a NAS Slammer. The Pilot manages data movement, but the actual movement occurs on a separate NAS Slammer.

NDMP Agent — This agent supports the Network Data Management Protocol (NDMP) that integrates third-party data management applications with system backup functions. It provides the control path to manage backup operations with local and network-attached tape libraries.

Configuration Server — The Pillar Axiom Configuration Server performs the following functions:

- Configures system storage and data paths.
- Manages system startup.
- Performs Slammer CU failover and failback.
- Manages drive replacement, including copyback and spare drive assignment.
- Manages RAID CU failures.
- Manages scheduled events.
- Manages Guided Maintenance.
- Monitors system configuration and status.
- Manages log collection, rotation, and transfer (automatic and manual).
- Manages and applies Quality of Service (QoS) policies.
- Provides the external interfaces to manage the system.

Configuration UI — The Pillar Axiom Configuration user interface (UI) component supports the set of user interfaces into the

system for management and control. The supported UIs are a Java UI and an installable command line application.

Pilot Persistence Pilot persistence stores software, logs, and configuration cache on the Pilot. Persistence is managed by the Pilot Configuration process.

Pilot Configuration The Pilot Configuration process manages all services on the Pilots, hardware initialization and control, logging, and Pilot persistence. It interacts with hardware and basic operating system functions to allow other code to function, including:

- Interprocess communication (IPC) and remote procedure call (RPC) mechanisms
- Memory allocation
- Hardware initialization and control
- Network protocol driver TCP/IP stack
- Topology discovery for storage control units
- Active or passive state management

SMIProvider The SMIProvider (Storage Management Initiative provider) is an SMIS-3 compliant Storage Networking Industry Association (SNIA) interface. SMIProvider can be used to monitor and manage storage on the Pillar Axiom system.

SNMP Agent The Simple Network Management Protocol (SNMP) agent provides a standard interface through the external management connection, which supports the SNMP protocol. SNMP GETs (queries) and Traps (event notifications) are supported.

Slammer Storage Controllers

The Slammer provides an external interface to the host storage network, and it processes every I/O request. A Pillar Axiom system can include both NAS and SAN Slammers.

Pillar Axiom systems support one, two, three, or four Slammers.

The following figure shows the different parts of a Slammer.

Figure 3 Slammer components

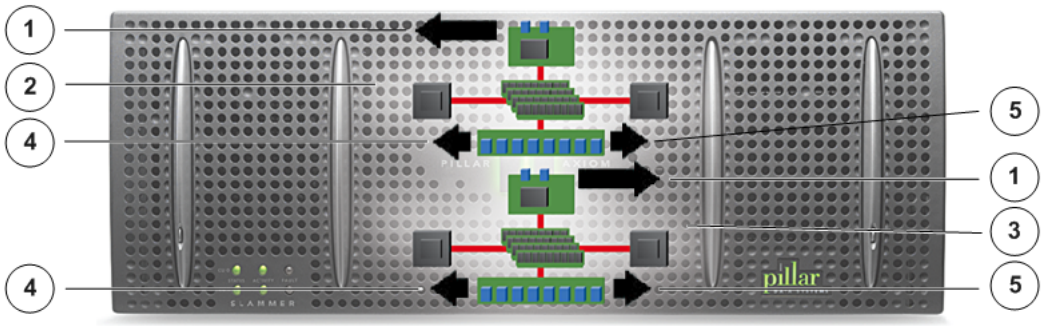

Legend	
	1 To host storage network (GbE, FC or iSCSI)
	2 Control unit CU0
	3 Control unit CU1
	4 To Bricks and other Slammer control units (storage system fabric)
	5 To Pilot (private management interface)

Slammer Functional Description

A Slammer contains two control units (CUs) functioning as an active-active asymmetric access pair. The primary components of each Slammer CU include:

- One private interface module (PIM).
- One network interface module (NIM) of one of the types described in Table 3: NIM configurations in each Slammer CU below.

- One dual-core or quad-core processor, as described in Table 4: Processors, memory, and power supplies in each Slammer CU below.
- Two power supplies, as described in Table 4: Processors, memory, and power supplies in each Slammer CU below.

The characteristics of these primary Slammer CU components are described in the following tables:

Table 2 PIM components in each Slammer CU

Components	Characteristics
Fibre Channel (FC) controller ports	Four, with one of the following switches: Version 1: 13-port FC switch with a 2 Gb/s copper back end Version 2: 16-port FC switch with one of: • 2 Gb/s copper back end • 4 Gb/s optical back end
Ports for the private management interface (PMI)	Three Ethernet ports
Status indicator on the Storage System Fabric (SSF)	One bi-color light-emitting diode (LED) for each port
Fault, status, and activity indicators	Two LEDs

Note: Because version 1 (13-port) and version 2 (16-port) Slammer PIMs use different components, version 1 and version 2 PIMs cannot co-exist in the same Slammer. However, a multi-Slammer system can contain a mix of version 1 and version 2 Slammers, as long as the PIMs within each Slammer use the same version. Also, because the presence of a 2 Gb/s component in an otherwise 4 Gb/s loop will force the loop to function at 2 Gb/s, 4 Gb/s optical and 2 Gb/s copper back ends cannot coexist effectively in the same Pillar Axiom system.

Table 3 NIM configurations in each Slammer CU

Slammer type	NIM Configuration
NAS Slammer	Four copper or optical ports with 1 Gb/s Ethernet RJ45 connectors (copper) or SFP transceivers (optical), or
	Two copper or optical ports with 10 Gb/s Ethernet SFP+ (Plus) transceivers
SAN Slammer, FC only	Two optical FC ports, each of which supports 1, 2, or 4 Gb/s FC SFP transceivers, or
	Two optical FC ports, each of which supports 2, 4, or 8 Gb/s FC SPF+ (Plus) transceivers
SAN Slammer, iSCSI only	Two copper iSCSI ports, each of which supports 1 Gb/s Ethernet SFPs
SAN Slammer, Combo FC and iSCSI	Two optical FC ports, each of which supports 1, 2, or 4 Gb/s FC SFPs, plus two copper iSCSI ports, each of which supports 1 Gb/s Ethernet SFPs, or
	Two optical FC ports, each of which supports 2, 4, or 8 Gb/s FC SPF+ (Plus) transceivers, plus two copper iSCSI ports, each of which supports 1 Gb/s Ethernet SFPs

Table 4 Processors, memory, and power supplies in each Slammer CU

Components	Series 1 Slammers	Series 2 Slammers	Series 3 Slammers
Processors in each control unit (CU)	One 2.6 GHz dual-core processor	One 2.2 GHz quad-core AMD Opteron family processor	One 2.2 GHz quad-core AMD Opteron family processor
Memory in each CU	12 GB	24 GB	24 GB
Power supplies in each CU	Two redundant power supplies	Two redundant power supplies	Two redundant power supplies

Both NAS and SAN Slammers are connected to the system storage pool through up to eight FC ports (four for each Slammer CU) on the NIM. On the PIM, Version 1 (legacy) ports are capable of 2 Gb/s connections only, and Version 2 ports are capable of 2 Gb/s or 4 Gb/s connections. The internal FC fabric is

connected through a set of FC loop switches. The Slammer virtualizes the storage pool, so you can easily increase the number and size of filesystems and LUNs.

Note: The Slammers do not perform RAID processing. RAID processing is handled at the Brick level.

Slammer Software Components

Most of the Pillar Axiom system software functionality resides in the Slammer. The software stack is logically separated by NAS and SAN functionality and management functionality. The Slammer software stack includes the following software components:

Array Manager	The Array Manager (AM) provides storage pool virtualization functions. It allows all data on Bricks to be treated as a single storage pool. The AM handles LUN mapping to physical drive storage locations.
Block Cache	The Block Cache component supplies all read and write cache management. All write operations are mirrored to the redundant Slammer control unit (CU) cache, providing full data integrity through a range of restart and failure scenarios. In addition, all snapshot actions are processed by the Block Cache service.
Common Internet File System (CIFS) Protocol	The CIFS protocol provides Windows and other CIFS clients access to a Pillar Axiom filesystem through the Ethernet ports. Storage is presented as CIFS shares.
Configuration Manager	The Configuration Manager provides the main management interface for system Slammer and Brick storage resources. Slammer software components get information about storage resource configuration and status from Configuration Manager, and they use Configuration Manager to change resource configuration. Configuration Manager provides a high-level system resource management interface that handles command distribution to appropriate Slammers and Bricks and manages the specific roles of other Slammer and Brick software components in carrying out requested operations. Specific functions Configuration Manager performs include:

- Create, modify, and delete volumes (filesystems or LUNs).
- Manage volume growth and migration.
- Manage the persistence of volume attributes and system settings.
- Coordinate initialization, de-initialization, and reconfiguration of system components during system startup, shutdown, and failure, as necessary.
- Implement high-level monitoring and management of Slammers and Bricks.
- Provide interfaces for querying the status of system Slammer and Brick resources, and of user-configured volumes.
- Manage partitioning of Brick storage into different Storage Domains.

Data Mover

The Data Mover provides the functions necessary for backup services. An API enables control from the Pillar Axiom Pilot Network Data Management Protocol (NDMP) daemon. The Data Mover communicates with the filesystem directly for both block-level and file-level transfers.

The Data Mover supports the following types of tape library attachments:

- Local attachment using Fibre Channel.
- Network attachment using Gigabit Ethernet. Supports both three-way and remote attachment.

Distributed Services Component

The distributed services component (DSC) provides a common set of services among all Slammer control units (CUs) in the system.

The DSC facilitates NAS and SAN integration. For example, this component makes it possible for a NAS Slammer to take a SAN LUN offline, or for a SAN Slammer to take a NAS filesystem offline. The DSC also allows a SAN Slammer CU to expand storage for a NAS Slammer CU.

Fabric Manager

The Fabric Manager manages all pathways within the SSF. Through the Fabric Manager, system components can send

	requests to each other without actual knowledge of the channel used to service the request. The Fabric Manager handles path failures by routing requests through alternate channels within the fabric.
Generic Data Mover	The Generic Data Mover is a platform-supplied service that provides data transfer and messaging requests across the Storage System Fabric (SSF).
InterConnect	The InterConnect communicates directly with the FC circuitry.
Meta Filesystem	The Meta Filesystem provides a protocol-neutral file system supporting files, directories, and other filesystem objects. It uses a transaction journal and advanced read caching algorithms to maximize I/O throughput. All filesystem I/Os are processed through the Virtual LUN (VLUN) layer.
Network File System (NFS) Protocol	The NFS protocol provides UNIX, Linux, and other NFS clients access to a Pillar Axiom filesystem through the Ethernet ports. Storage is presented as NFS mount points.
Platform Services	Slammer Platform Services provides the code necessary for the Slammer to interact with hardware and basic system operational functionality. Specific services include: ○ Hardware initialization and control ○ Interprocess communication (IPC) and remote procedure calls (RPC) ○ Memory allocation ○ Network protocol drivers ○ Small Conputer System Interface (SCSI) protocol drivers ○ Fibre Channel (FC) drivers for private interconnect ○ RS-232 support ○ Private interconnect switch management ○ Diagnostic management ○ Topology discovery
SCSI Command and Control	This layer includes host FC and iSCSI drivers as well as the command processing function for SAN attachment. The

SCSI Command and Control component processes all I/Os through the VLUN layer.

Brick Storage Enclosures

There are three types of Bricks:

- SSD (solid-state drive) Bricks
- SATA (serial advanced technology attachment) Bricks
- FC (Fibre Channel) Bricks

Each of these Brick types is a unique Storage Class. All three Brick types can be mixed in a single Pillar Axiom system. All three Brick types can also be mixed in a single Brick chain, but we recommend separating Brick types into distinct chains, if possible.

All types of Bricks support both RAID 5 and Distributed RAID within the same array of the Brick simultaneously.

The drives used in SSD Bricks feature:

- SATA interface
- SSD media
- 50 GB or 200 GB capacity

The drives used in SATA Bricks feature:

- Multiplexed SATA interface
- Hard disk drive (HDD) media
- 400 GB (legacy), 500 GB, 750 GB (legacy), 1 TB, or 2 TB capacity

The drives used in FC Bricks feature:

- Fibre Channel Interface
- HDD media
- 146 GB (legacy), 300 GB, 450 GB, or 600 GB capacity
- Concurrent dual-port transfers

Brick Functional Description

The architecture of a Brick provides built-in redundancy of its components. Each RAID controller on a serial advanced technology attachment (SATA) Brick acts

as an active controller for one of the two RAID arrays in the Brick and has paths to each of the drives inside the Brick. On Fibre Channel (FC) Bricks, both controllers can access all the drives in the Brick. If one controller on a Brick fails, the other controller continues to process I/Os for all arrays within the Brick. The RAID controllers are hot-swappable.

Data stored in a Brick is accessed by the Slammers through the Fibre Channel-based Storage System Fabric (SSF). A Slammer can read or write data to or from any RAID array on any Brick, because storage resources are shared across all controllers. Standard SCSI commands enable communication between Slammers and Bricks.

Advantages of the Brick RAID architecture include:

- Rebuilding from a failed drive is managed at the Brick level, minimizing the performance impact during the rebuild and drastically shortening the rebuild time.
- Because RAID processing is performed at the Brick level, not the Slammer level, Brick RAID controllers free up the Slammer to focus on I/O instead of RAID processing. This improves scalability because the number of I/O operations each second (IOPs) increases as capacity is added.
- If one RAID controller fails, the remaining RAID controller has enough processing power to handle the entire array, so there is no degradation of performance.
- RAID controllers have redundant Fibre Channel cross connections to increase the number of paths from any Slammer control unit (CU) to any RAID array.
- Drives are mounted horizontally in the tray to facilitate drive replacement.
- The Brick firmware can perform enhanced drive error recovery by temporarily removing a suspect drive from the storage array and performing extended recovery actions on that drive. This promotes drive reliability.

  The Brick firmware monitors SMART (self-monitoring, analysis, and reporting technology) data for each individual drive. The firmware will proactively remove a drive from service if the drive exceeds SMART error thresholds and there is a spare drive available.
- Redundant power and cooling ensures that a Brick will continue to function normally with only one power supply active.

The controllers monitor the following Brick status information:

- Internal temperatures

- Fan speed control and feedback
- Drive carrier detection
- World wide name (WWN) and system serial number information
- Power supply status and detection information

The RAID controller tracks system resets. The RAID controller can also track the removal or insertion of a Fibre Channel cable.

SSD Bricks

Solid-state drive (SSD) Brick enclosures contain 13 SSDs, of which 12 are arranged in two six-drive arrays. The 13th drive is used as a hot spare for automatic failover.

SSD Brick storage enclosures are managed by a pair of version 2 RAID controllers.

Figure 4 SSD Brick components

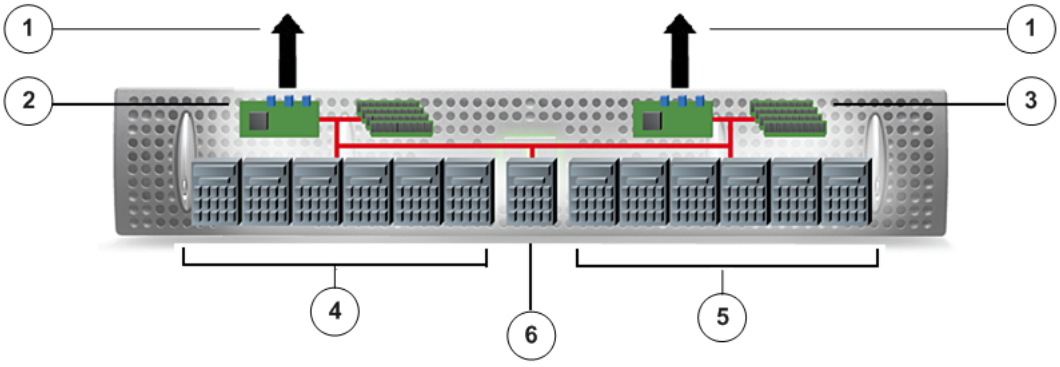

Legend	
	1 Connects to a Slammer or another RAID controller
	2 RAID controller
	3 RAID controller
	4 RAID group (six SSDs)
	5 RAID group (six SSDs)
	6 Hot spare

SSD Bricks use the same SATA interface as SATA Bricks.

The number of SSD Bricks supported depends on the number of Slammers:

- Single-Slammer systems support up to eight SSD Bricks.
- Two and three-Slammer systems support up to 16 SSD Bricks.
- Four-Slammer systems support up to 32 SSD Bricks.

SATA Bricks

SATA Brick storage enclosures contain 13 hard disk drives (HDDs), of which 12 are arranged in two six-drive arrays. The 13th drive is used as a hot spare for automatic failover.

SATA Brick storage enclosures are managed by a pair of RAID controllers.

Figure 5 SATA Brick components

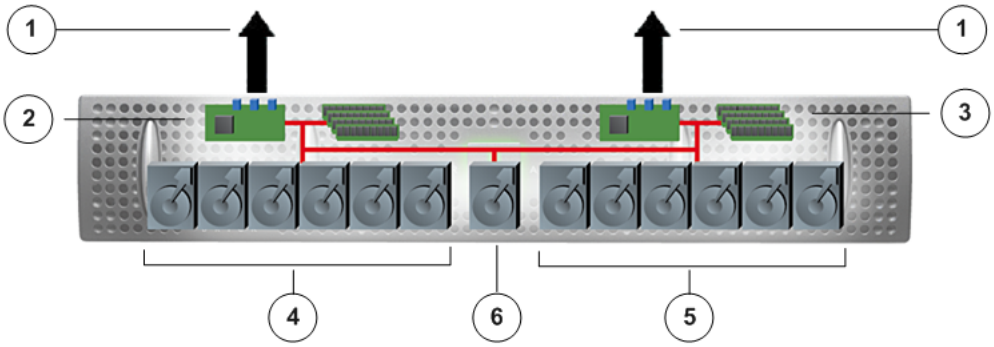

Legend	
	1 Connects to a Slammer or another RAID controller
	2 RAID controller
	3 RAID controller
	4 RAID group (six HDDs)
	5 RAID group (six HDDs)
	6 Hot spare

Under normal conditions, each controller provides access to and control over an array of six HDDs. Under failover conditions, a single controller can control and provide access to both arrays. All Pillar Axiom systems support SATA RAID controllers.

SATA Brick controllers come in two types:

- Version 1 (legacy) controllers have one set of four Fibre Channel (FC) ports with high speed serial data connector (HSSDC) connectors, and support only 2 Gb/s copper connections.
- Version 2 controllers have two pairs of FC ports with small form factor pluggable (SFP) connectors, and support either 2 Gb/s copper or 4 Gb/s optical connections. In addition, version 2 SATA controllers employ an updated chipset with greater internal bandwidth, and they support SATA HDDs as well as solid-state drives (SSDs).

Because version 1 and version 2 SATA controllers use different internal communication protocols, these two types of SATA controllers cannot co-exist in the same Brick chassis. In other words, you cannot use a version 2 SATA controller to replace a legacy version 1 controller. A Pillar Axiom system can, however, contain a mix of version 1 and version 2 SATA Bricks.

For a complete list of the rules for configuring SATA Bricks, refer to the *Pillar Axiom SSF Cabling Reference* for the Pillar Axiom system version being configured.

The number of SATA Bricks supported depends on the number of Slammers:

- Single-Slammer Pillar Axiom systems support up to 32 SATA Bricks.
- Two, three, and four-Slammer Pillar Axiom systems support up to 64 SATA Bricks.

Fibre Channel Bricks

Fibre Channel (FC) Brick storage enclosures contain 12 hard disk drives (HDDs) arranged in a single 11-drive array plus a hot spare. FC Bricks do not have a dedicated hot spare; instead, any drive can be utilized as a spare. If a drive fails, the rebuild occurs on the current hot spare. After the failed drive has been replaced, it becomes the new hot spare.

FC Brick storage enclosures are managed by a pair of RAID controllers.

Figure 6 FC Brick components

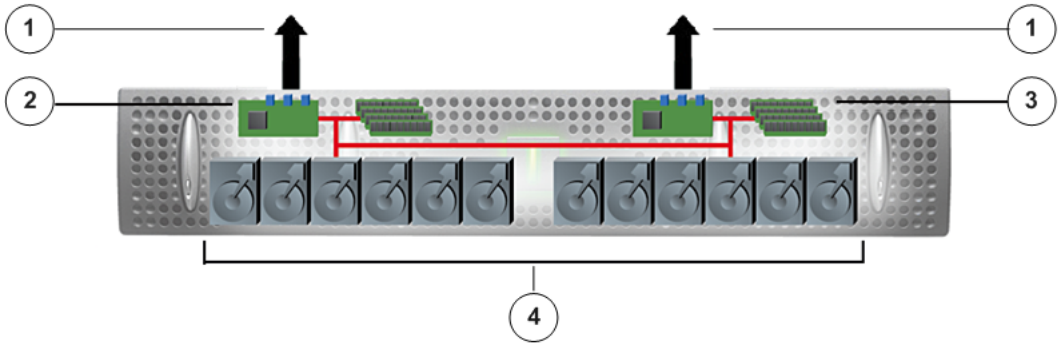

Legend	
	1 Connects to a Slammer or another RAID controller
	2 RAID controller
	3 RAID controller
	4 RAID group (12 HDDs including a hot spare)

FC Brick controllers come in two types:

- Version 1 (legacy) controllers have one set of four Fibre Channel (FC) ports with high speed serial data connector (HSSDC) connectors, and support only 2 Gb/s copper connections.
- Version 2 controllers have two pairs of FC ports with small form factor pluggable (SFP) connectors, and support either 2 Gb/s copper or 4 Gb/s optical connections. In addition, version 2 FC controllers employ an updated chipset with greater internal bandwidth, and they support FC HDDs as well as solid-state drives (SSDs).

Because version 1 and version 2 FC controllers use different internal communication protocols, these two types of FC controllers cannot co-exist in the same Brick chassis. In other words, you cannot use a version 2 FC controller to replace a legacy version 1 controller. A Pillar Axiom system can, however, contain a mix of version 1 and version 2 FC Bricks.

Unlike other Bricks, the number of FC Bricks supported does not depend on the number of Slammers. Single, two, three, and four-Slammer Pillar Axiom systems support up to 32 FC Bricks.

SSD, SATA, and FC Bricks can co-exist on the same Brick string subject to configuration recommendations.

A given Brick string can contain up to a total of four FC Bricks.

For a complete list of the rules for configuring FC Bricks, refer to the *Pillar Axiom SSF Cabling Reference*.

Brick Software Components

The following software resides in the Brick:

Target Fibre Channel Driver

The Target Fibre Channel (FC) driver serves as an isolation layer between the FC hardware and the remainder of the Brick firmware. This allows the FC hardware to change without great impact upon the remaining code. The FC driver also translates hardware-specific sets of data structures to simple Small Computer System Interface (SCSI) requests that are queued and eventually processed by the SCSI layer of the Brick firmware.

The FC driver includes a transfer manager that facilitates data transfers from the Brick's data buffer across the FC link to and from the requestor. This includes translating scatter/gather lists into the format expected by the FC hardware.

SCSI Layer

The SCSI layer receives queued SCSI requests from the target FC driver. It validates these requests and converts them into the standardized command request format used by the rest of the Brick firmware. This conversion includes correcting command parameters from the big-endian data format commonly used in SCSI requests to the native data format of the processor used by the Brick. Once the request is parsed, the SCSI layer sends the newly created command request block on to the command processor.

Command Processor

The command processor accepts command request blocks from the SCSI layer and dispatches these requests to the various routines responsible for their timely completion.

Cache Management

The cache management code supplies buffer memory for use in processing commands. The cache manager also maintains previously used valid data as long as possible so that this cached data can be used to satisfy future requests without requiring access to slower media types. Should the data needed for a request not be available in cache, the cache code issues the RAID requests needed to gather the data from attached drives.

RAID Engine	The RAID engine converts RAID requests into one or more drive requests directed to the storage media. These drive requests are queued and eventually processed by the drive manager. This process requires the RAID engine to calculate the logical-to-physical mapping for various RAID types so that a given RAID unit address always accesses the same physical media address.
Partition Manager	The partition manager creates, saves, and restores the information required to define the logical RAID units used by the RAID engine.
Drive Manager	The drive manager is responsible for handling all drive requests and any error recovery required to complete those requests. It also maintains the current state for all physical media (drives) in the Brick, and updates that state based upon the results of each I/O operation.
Storage I/O Driver	The storage I/O driver provides an interface to the initiator hardware that the FC or SATA driver provides for the target hardware. The storage I/O driver converts the drive requests issued by the RAID engine into the hardware-specific request structures that the initiator hardware requires to communicate with the storage media. This driver also acts as an isolation layer between most of the Brick firmware and the initiator hardware, minimizing code changes when hardware is updated.
Pre-Emptive Copy	Pillar Axiom Pre-Emptive Copy initiates recovery of a drive before the drive fails. By simply copying the data to the spare drive rather than rebuilding the failing drive, pre-emptive copy shortens recovery time and reduces the chance of data loss.

CHAPTER 3

Pillar Axiom Software Overview

Pillar Axiom Storage Services Manager

The Pillar Axiom Storage Services Manager is an easy-to-use GUI. It is organized into sections to help you configure and monitor your Pillar Axiom system. These sections appear at the top of the left navigation pane in the following figure.

Figure 7 Pillar Axiom Storage Services Manager

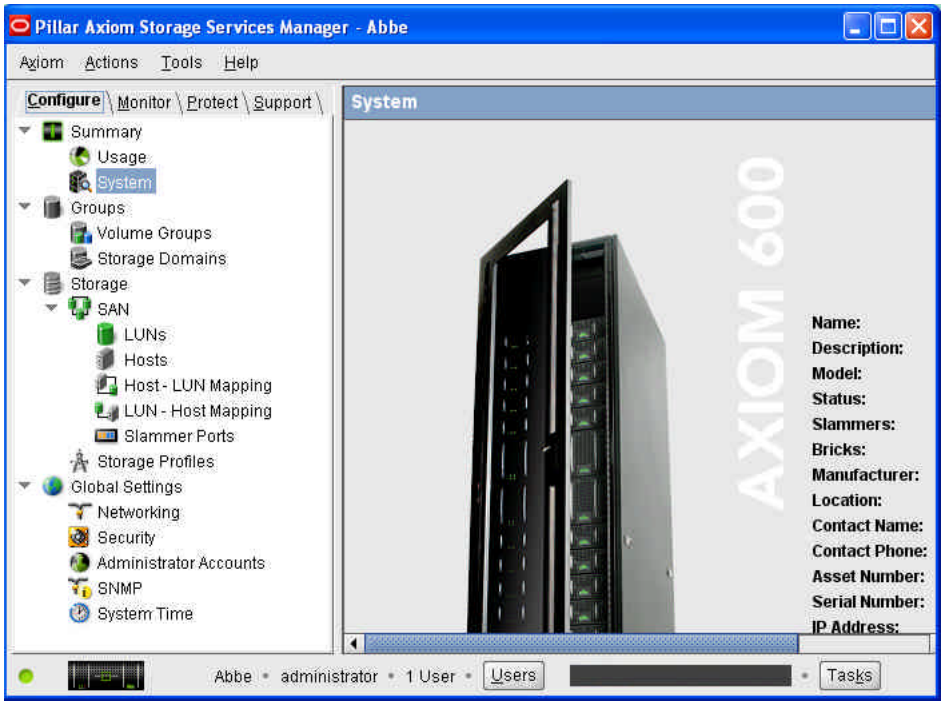

These sections perform the following functions:

- **Configure**: Sets up global settings, such as networking, security, and administrator accounts. Creates and manages storage objects such as volumes (LUNs and filesystems), Storage Domains, and storage profiles.
- **Monitor**: Keeps track of system hardware, system alerts, and event notifications. Provides access to Guided Maintenance, which enables storage administrators to identify and replace failed components, often without interrupting system operation. Also manages scheduled tasks (such

as data protection schedules), generates statistical reports, and displays tables and charts of system statistics.

- **Protect**: Schedules backup operations and provides immediate data replication.
- **Support**: Manages configuration of software modules, including scheduled software updates. Provides tools for monitoring system logs, and troubleshooting. Also provides access to technical documentation, utilities downloads, and contact information.

The status bar at the bottom of the Pillar Axiom Storage Services Manager screen provides instant feedback on system performance, running background tasks, or system alerts that require attention. For detailed information, refer to the *Pillar Axiom Administrator's Guide*.

Policy-Based Provisioning and Storage Management

The Pillar Axiom system manages storage resources using administrator-defined policies, which are the basis of the storage management system. Performance, utilization, and availability metrics are tailored to individual logical volumes. Policies are established through the Pillar Axiom user interface, using the graphical user interface (GUI) or the command line interface (CLI), and implemented in the core layer of the Pilot. The Configuration Manager facilitates the control, monitoring, and reporting of the Slammers to help the Pilot enforce these policies.

The Pillar Axiom system can create a filesystem or LUN to match data performance, relative priority, and access pattern. Standard or compliance retention policies can also be applied to filesystems. This flexibility ensures that applications ranging from mission-critical to archive receive the appropriate system resources.

The Pillar Axiom Storage Services Manager can also modify the performance, priority, access, or retention properties of existing volumes. If necessary, this process can automatically migrate user data to a different physical location on the storage pool to fulfill the modification request.

Quality of Service Attributes

The Pillar Axiom system defines a set of policies, or Quality of Service (QoS) attributes, that govern the QoS for the volume (filesystem or LUN). These policies determine how data is stored in the storage pool, the queueing and caching priority of the data, and the type of RAID used for the data, so that the highest performance is achieved.

A Pillar Axiom system allocates storage with different application priorities. System resources are disproportionately applied to deliver the requested QoS for each volume. Storage administrators can allocate storage resources and define storage automation parameters through the Pillar Axiom Storage Services Manager, and the software takes care of the rest.

A single Pillar Axiom system can deliver many tiers of storage, differentiated by performance, availability, data protection, capacity, and scalability. A common implementation has three tiers of storage services:

- Tier 1: Highest performance and availability levels for mission-critical applications, represented as red in the following figure.

- Tier 2: Increased performance and high availability for mid-performance applications, represented as yellow in the following figure.
- Tier 3: Adequate performance and availability to support business-utility applications such as file sharing and archive, represented as blue in the following figure.

Figure 8 Pillar Axiom system data layout

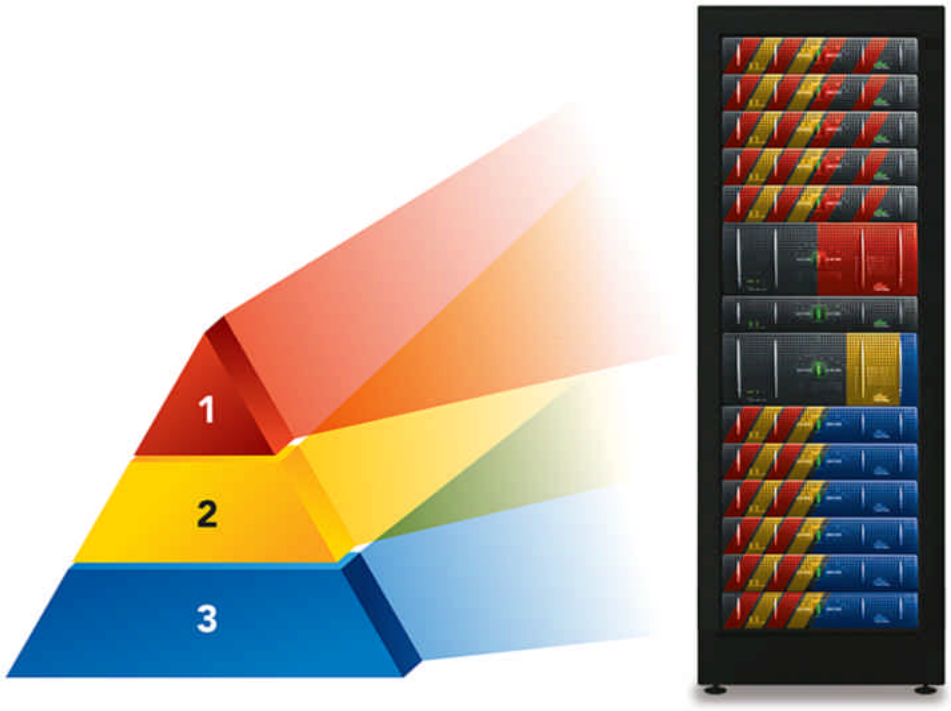

The Pillar Axiom Storage Services Manager makes it possible to assign array resources for each volume (filesystem or LUN), much the same as configuring a virtualized server resource.

Dynamic provisioning and profile-based resource assignments provide the necessary flexibility for quickly and easily adjusting capacities to meet ever changing business storage demands. All of the QoS settings, along with the redundancy attributes, are utilized to determine the RAID type of a volume (filesystem or LUN).

Storage Profiles

When configuring a logical volume, you can select a collection of predefined properties to apply to that volume. This collection of properties is called a *Storage Profile*.

When using a specific Storage Profile, you can select a profile that you have previously created and saved or one of the pre-configured profiles.

After a volume is created using a Storage Profile, removal of the profile does not affect the performance characteristics of that volume.

Storage Classes

The Storage Class feature allows you to specify the preferred storage media to use for a logical volume.

A Storage Class is defined as:

> A categorization of physical storage, each category having distinct characteristics with regard to performance characteristics of data access. Example Storage Classes in a Pillar Axiom system are serial ATA (SATA), Fibre Channel (FC), and solid state drive (SSD). Pillar Axiom systems allow an administrator to explicitly manage volume placement within the overall system storage pool, first by Storage Domain, then by Storage Class, and finally by relative priority level within that Storage Class.

Pillar Axiom systems support the following three Storage Classes:

- SATA
- FC
- SSD SLC (solid state drive, single-level cell)

Note: Which Storage Classes are available on a particular Pillar Axiom system depends on the types of Brick storage enclosures you have installed on the system.

A Storage Class has these attributes:

- A newly created logical volume is associated with a single Storage Class.
- The Pillar Axiom Storage Services Manager graphical user interface (GUI) shows the capacity available within each Storage Class.
- The system will not create a logical volume when the available space for the associated Storage Class is insufficient to accommodate the capacity requested for the volume.

For FC and SATA Storage Classes, the striping of a logical volume is across a number of drives in a collection of RAID groups. The number of drives depends on the Quality of Service (QoS) priority setting for the volume. For the SSD SLC

Storage Class, striping for a volume is across all available drives, regardless of the priority setting.

Storage Domains

Storage Domains allow storage administrators to assign logical volumes to a specific collection of Bricks. Such assignments can be made to reduce contention among volumes, to implement different levels of security for those volumes, or both.

Note: Storage Domains might limit the ability of the system to provide the best optimization of the storage arrays and system performance.

A Storage Domain is defined as:

> A subset of a virtual storage pool consisting of a defined group of Brick storage enclosures. This group can consist of any assortment of Bricks, regardless of Storage Class, capacity, or any other attribute. A Storage Domain is typically used to provide specific allocation or security features for a collection of logical volumes.

An administrator can allocate each Brick to a defined Storage Domain. When no administrator-defined domains exist, all Bricks reside in the default domain.

Storage administrators typically use Storage Domains for the following reasons:

User group separation	In this scenario, storage administrators can isolate application data to specific Bricks on a department basis (for internal cloud environments) or on a customer basis (in external cloud environments). This isolation eliminates inter-application contention for I/O services and provides charge-back capabilities.
Protocol separation	In this scenario, storage administrators can place application data on separate Bricks based on protocol and connectivity. This separation eliminates any chance of inter-application contention for I/O services. For example, an administrator could create a NAS domain, a SAN iSCSI domain, and a SAN FC domain.

Application I/O isolation	Storage administrators can create Storage Domains for use in specific applications and tiers of storage to eliminate unwanted Brick contention. For example, an administrator can create a replication domain for incoming replication data and another domain for archival or backup of local data.
Data security	Storage administrators can place logical volumes that contain sensitive data on a particular Storage Domain. If the data needs to be destroyed, the drives within those Bricks can be destroyed without the administrator having to be concerned with preserving less sensitive data. Placing those volumes in their own Storage Domain ensures that those volumes do not share Bricks with less sensitive material.
Brick or hardware retirement	As drives age, the probability of failure increases. Storage Domains can efficiently move data to newer Bricks that have larger capacities as well as updated RAID controllers.

Figure 1 illustrates a collection of Storage Domains and a sample distribution of logical volumes across those domains. This illustration shows the relationships among the following collection of objects:

- Three Storage Domains
- Two volume groups (one nested)
- Five logical volumes
- Seven Bricks

Figure 9 Storage Domains, volume groups, and volumes

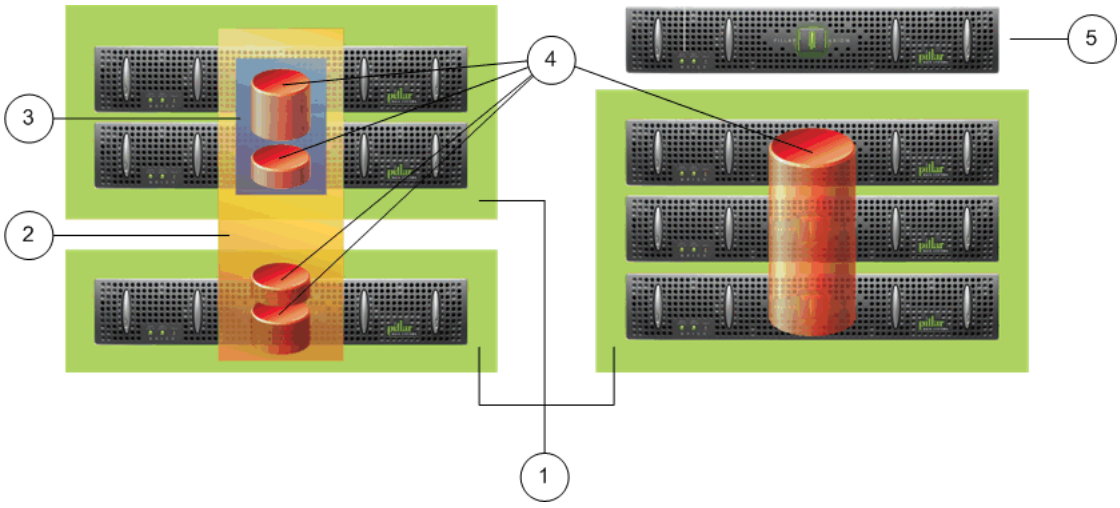

Legend		
	1 Storage Domains	4 Logical volumes
	2 Volume group	5 Unassigned Brick
	3 Volume group (nested)	

In the illustration, the outer volume group (item 2, the orange box) contains a nested volume group (item 3, the blue box). The nested volume group contains two logical volumes (item 4, the red cylinders), while the outer (or parent) volume group contains two volumes of its own. Volume groups can also span multiple Storage Domains.

Note: Volume groups are always optional, as illustrated by the Storage Domain on the right side of the illustration, which contains a volume that is not part of a volume group.

The preceding figure also shows an example of a Brick that is not assigned to any Storage Domain. This state is temporary. While in this state, the capacity of the Brick is not included as free or available capacity. Causes of an unassigned state for a Brick:

- Newly added to the system
- About to be removed from the system
- In transition from one Storage Domain to another

Storage administrators can perform regular management actions for any logical volume residing in a Storage Domain, including:

- Create logical volumes within a domain.
- Create Volume Copies within a domain.

- Create clones of logical volumes contained in a domain.
- Move logical volumes to a different volume group.
- Delete logical volumes from a domain.

Note: All allocation for a logical volume is confined to the Bricks within a Storage Domain. In other words, the extents associated with a volume cannot span more than one domain.

Primary Storage Domains

Each Pillar Axiom system has exactly one primary Storage Domain. This domain contains system overhead, including all system configuration data.

RAID Array Stripes

Pillar Axiom systems support RAID 5 and Distributed RAID geometries within the same Brick array.

Strips are disk block addresses. A RAID array stripe consists of a set of consecutively addressed strips.

RAID 5 arrays support the following strip sizes:

- For wide stripes: 1 MB for each strip.
- For standard stripes:
  - Fibre Channel (FC) Bricks: 64 KB for each strip.
  - Serial ATA (SATA) and solid-state drive (SSD) Bricks: 128 KB for each strip.

Distributed RAID arrays are formed from pairs of standard strips (64 KB strips for FC and 128 KB strips for SATA and SSD) only.

For FC Bricks, a stripe is a collection of 10 data strips and one parity strip. Each strip (64 KB) is written to one of the drives in a FC Brick, which means the stripe is written across 11 drives. For FC Bricks, a stripe also contains 640 KB, but its width is 11. Each FC Brick contains one such array, plus a hot spare.

For SATA and SSD Bricks, a stripe is a collection of five data strips and one parity strip. Each strip (128 KB) is written to one of the drives in a RAID array, which means the stripe is written across six drives. For SATA and SSD Bricks, a

stripe contains 640 KB, and its width is six. Each Brick contains two such arrays, plus a hot spare.

For an Oracle Automatic Storage Management (ASM) storage profile, strips contain 1024 KB (1 MB). The number of strips for each stripe remains the same, depending on the type of Brick. Also, the stripe width does not change, only the size of the strip does.

Thinly Provisioned Volumes

The Pillar Axiom system allows you to provide thinly provisioned volumes (filesystems and LUNs).

Thin provisioning is defined as:

An approach to storage allocation in which a logical volume appears to be much larger than the storage actually allocated to it. Additional storage is dynamically allocated when necessary. Administrators interact with thinly provisioned volumes when configuring their capacity and growth increments. These types of volumes are sometimes referred to as *sparse filesystems* and *sparse LUNs*.

Note: A Pillar Axiom system uses binary units to calculate and display the capacity of physical storage and the size of logical volumes:

1 MB = 1024^2 (1,048,576) bytes

1 GB = 1024^3 (1,073,741,824) bytes

1 TB = 1024^4 (1,099,511,627,776) bytes

The following sections describe thin provisioning and how it affects storage capacity.

About Thinly Provisioned Volumes

Traditionally, when storage is allocated to an application, the allocation is dedicated to that application. This assignment prevents other applications from accessing this capacity, even when the amount allocated is never used. Because of this allocation strategy, the capacity is stranded and cannot be leveraged in support of additional needs.

Thin provisioning mitigates these issues by allowing storage administrators to leverage this unused capacity for a logical volume by performing these actions:

- Allocate capacity based on future needs.
- Draw on a common pool of storage as capacity is consumed.

Thin provisioning allows an administrator to create a logical volume of any size without committing that capacity at that time. Each application has what appears to be all the storage needed for ongoing operations, but without the physical capacity locked to a particular volume.

Administrators can create logical volumes up to the maximum addressable logical capacity that is allowed for the OS, with little physical storage assigned to the volume. As data is written to the thinly provisioned volume and capacity is consumed (called *in-fill*), the system automatically allocates additional capacity to the logical volume in increments.

Note: Solid-state drive (SSD) Bricks do not support thinly provisioned volumes.

Free Capacity in Thinly Provisioned Volumes

A minimum amount of free space is required to create a new logical volume. The actual amount of physical capacity that is consumed from the system free space when you create a new logical volume depends on several factors.

These factors are:

- The RAID geometry of the volume.
- The redundancy Quality of Service (QoS) setting of the volume.

To determine the actual physical capacity needed, the system adds the following:

- To account for parity, the system increases the requested capacity by different amounts, depending on the RAID geometry:
  - 20% for RAID 5 (SATA)
  - 10% for RAID 5 (FC)
  - 100% for Distributed RAID or RAID 5 with Wide Stripe
- If redundancy for the volume is set to Double, the system doubles the physical allocation.

For example, if the requested capacity for a logical volume is 250 GB, and the volume uses RAID 5 geometry in SATA storage, the system allocates an additional 50 GB. If the volume has a redundancy setting of Double, the system allocates an additional 300 GB, for a total physical allocation of 600 GB.

Storage Allocation of Thinly Provisioned Volumes

The capacity reserved for thin provisioning, which is part of the system overhead, is accounted for in the available capacity that the system reports. In other words, what the system reports as available capacity is fully available for the provisioning of logical volumes.

Unused capacity in the storage array can decrease over time. This decrease is due primarily to two events:

- New volumes are created.
- Thinly provisioned volumes are provisioned (filled in) when they grow. When no unused system capacity remains, the system uses this reserve to fill in the thinly provisioned volumes.

For storage area network (SAN) systems, the degree to which a LUN is thinly provisioned depends on the nature of the host applications that access the LUN. If only specific portions of a LUN are ever accessed by applications, the thinness of that LUN remains the same. As applications attempt to access more and more different areas of the LUN, the system allocates more and more physical space for the LUN, causing the thinness to decrease.

For network attached storage (NAS) systems, the degree to which a filesystem is thinly provisioned depends on the maximum amount of space ever used by this filesystem. As a filesystem consumes more space, it requires more allocation of physical storage to become less thin.

Reducing the space used by a filesystem (by deleting files or snapshots, for example) will not result in physical storage being freed. Thus, reducing the space used by a filesystem will not increase the thinness of the filesystem.

Growth Increments

When the system allocates capacity for a logical volume, the system divides the allocation into slices (called *growth increments*) and uses as many of them as it needs.

Each growth increment is between 1 and 2 GB. For example, if the volume is 2 TB, the system may use between 1024 and 2048 growth increments for the allocation. The exact value depends on the combination of the following choices that characterize the underlying storage for the volume:

- Type of Brick (SSD, Fibre Channel, or serial ATA)

- RAID geometry (RAID 5 or Distributed RAID)
- Strip size (normal or 1 MB)

Note: When the system needs to grow or in-fill a logical volume, the system returns an error if sufficient capacity does not exist within the Storage Class associated with the volume, even when sufficient capacity exists in other Storage Classes.

Capacity Overhead

Plans for the provisioning of logical volumes must take into account the extra capacity the system allocates to overhead.

To accommodate the level of RAID protection required to allocate a newly created logical volume, the system adds a certain amount of overhead to a request for the capacity of the volume. The capacity consumed and reported for RAID 5 logical volumes includes that overhead. This overhead varies, depending on the RAID geometry and Storage Class assigned to the volume. For RAID 5, the overhead is as follows:

Serial ATA drives and SSDs	20%
Fibre Channel drives	10%

For Distributed RAID, the capacity consumed and reported for logical volumes is twice the requested amount, regardless of Storage Class.

Besides the overhead allocated to a logical volume when the volume is created, the Pillar Axiom system allocates 50 GB of physical capacity in each of the serial ATA (SATA) and Fibre Channel (FC) Storage Classes as an in-fill reserve. The system reserves this physical capacity to help prevent inadvertent exhaustion of system physical capacity when thinly provisioned volumes are created. The system uses this capacity when physical capacity needs to be assigned to a thinly provisioned volume, and all other physical capacity in that Storage Class has been consumed.

The size of this reserve capacity is included in the calculations for the free, available, and total system capacities that are displayed by the graphical user interface (GUI) and the command line interface (CLI).

Parity in Reported Capacities

RAID arrays have both physical and virtual capacity.

The physical capacity of a RAID array that is reported includes capacity for parity. Sizes reported in capacity usage summaries and the sizes reported for total, used, and free system capacities are in terms of raw physical capacities.

The virtual capacity of a RAID array that is reported, however, does not include capacity for parity. The ratio between the virtual capacity and the physical capacity depends on whether the storage is RAID 5 or Distributed RAID:

RAID 5: serial ATA (SATA) drives and solid state drives (SSDs)	5:6
RAID 5: Fibre Channel (FC) drives	10:11
Distributed RAID: FC, SATA, and SSD drives	1:2

Reclaiming Capacity

When a user deletes a logical volume, the system reconditions the space (by writing a predefined bit pattern) before reclaiming it for reuse. As the previously allocated capacity frees up, it becomes available for allocation.

Note: When a large volume is being deleted, the operation can take awhile for all the capacity to be reclaimed. Because of this additional time needed for reconditioning, the amount of used capacity plus the free capacity may not equal the total capacity. During this time, the graphical user interface (GUI) displays the amount of capacity remaining to be reconditioned.

For filesystems, when a user deletes a file that has no snapshots associated with it, the freed blocks appear as free capacity in the snapshot repository that is associated with the parent filesystem. Utilization commands (such as `df` or `du` on UNIX systems) show this newly freed capacity.

If the deleted file has snapshots associated with it, the system preserves the blocks in the snapshot repository for that filesystem. In this case, the number of free blocks for the filesystem does not change. Utilization commands show no change in the used and free space for that filesystem. To return these blocks to the free system space, all snapshots in the filesystem must be deleted.

Note: If, however, this deleted file was modified after the most recent snapshot and contained new blocks of data not captured by that snapshot, the system reclaims those new blocks. Utilization commands in this case would show those newly freed blocks as additional free capacity for that filesystem.

High Availability

Pillar Axiom systems are designed to provide the highest possible levels of availability.

Many Pillar Axiom components, such as network interfaces, control unit motherboards, power supplies, and fans, are redundant. This redundancy is intended to eliminate single points of failure.

The software layers of the Slammer support high availability. Software processes on each control unit (CU) are in constant communication with each other regarding their status and ability to perform.

The Pillar Axiom system uses a double-safe write system. The system secures the I/O in battery-backed, non-volatile RAM (NVRAM), so that the I/O is safe in case of external power loss. The Pillar Axiom system's redundant CU architecture secures the I/O in both CUs before the write operation is acknowledged as a complete transaction. The write is secured in two places, so only a catastrophic system event can affect the write integrity.

A double-safe write operation stores the data in local memory on the primary CU and in battery-backed memory on the alternate CU. If the primary CU fails completely, the alternate CU already has the data so it can recover and continue.

System Maintenance and Problem Detection

The Pillar Axiom system provides several different features to maintain the system and detect problems so that they can be resolved quickly. These features allow you to review system health, set up alert notification, review event log entries, generate reports, and display system statistics.

The following sections describe some of these features.

About Managing Call-Home Settings

Manages the Call-Home settings on a Pillar Axiom system and notifies the Pillar World Wide Customer Support Center of status and configuration information or any issues.

Call-Home is a feature that, when enabled, allows the system to send the status and configuration information to the Pillar World Wide Customer Support Center; no customer data is sent. The Call-Home feature also notifies the Pillar World Wide Customer Support Center about issues in the Pillar Axiom system. For example, when a component operates in degraded mode or fails, the system automatically performs failover actions. Although a component failure does not cause downtime, manual intervention is sometimes required to repair or replace the failed component. The system sends a Call-Home message to initiate the repair or replacement process.

Call-Home log collection can be initiated by one of the following methods:

- Manual: The administrator has requested a log collection.
- Event-triggered: An event has triggered the Call-Home.
- Periodic: A specified time has elapsed since the Call-Home was triggered.

The Pillar Axiom system maintains a directory of data files, each of which captures a Call-Home session. Whenever one of these data files is overwritten or thrown away, a log entry is made noting that fact. The collection of data files represent the ten most recent Call-Home sessions. The system administrator can select a session file and download it to a client machine or send it directly to the currently targeted server. Call-Home sessions can also be sent to a local Call-Home server. Contact the Pillar World Wide Customer Support Center for details.

Guided Maintenance

Guided Maintenance is a feature in the Pillar Axiom system that provides users with a method to identify and replace a field replaceable unit (FRU) in many cases without interrupting system operation and with minimal system degradation. There are four parts to Guided Maintenance:

- First, the system determines the status of the hardware based on events and diagnostics to accurately reflect the state of the system.
- Second, the system helps the administrator correctly identify the faulty component by presenting images of the FRU or beaconing the FRU.
- Third, if required, the system places system components in a condition to prepare for the replacement by redirecting activity from the faulty component to a redundant component.
- Fourth, the system guides the administrator through a set of procedures to replace the FRU.

Providing accurate system status and replacing FRUs are complex operations that involve many lower level components within the system. The process and details on how to maintain system operation are hidden from the user. The Pillar Axiom system is designed to be maintained by the user without requiring support from the Pillar World Wide Customer Support Center.

Each FRU has its own diagnostics which are called by the Pilot to verify that a FRU is accessible and functioning properly. The diagnostics are primarily used to verify FRUs that have been added or replaced. The system also tracks parts that have failed and been removed to prevent re-insertion of failed components.

Guided Maintenance supports the identification and replacement of FRUs for Slammers and Bricks. The *Pillar Axiom Service Guide* provides instructions on all replacement procedures.

Pillar Axiom Pre-Emptive Copy

The Pre-Emptive Copy feature further shortens the time it takes to rebuild a failed drive by doing the bulk of the work before a failure occurs, using a simple copy instead of a rebuild of the entire contents of the drive.

Pillar Axiom Pre-Emptive Copy is a Reliability, Availablity, Serviceability (RAS) feature of the Pillar Axiom RAID firmware that copies the data on a drive (which has been predicted to fail) to the spare drive before the suspect drive fails and is

subsequently taken offline for replacement. This RAS feature avoids performance degradation and potential exposure to data loss when the drive does fail.

Pre-Emptive Copy leverages S.M.A.R.T (Self-Monitoring, Analysis, and Reporting Technology) and other heuristics that provide detailed drive health status. It uses adaptive learning to discover any drive that might be in danger of failing and, before calling for a replacement, begins copying the data on that drive to the spare drive. Only the drive data (not the entire drive image) is copied, and the copy takes place in the background to ensure maximum system performance.

Software Updates

Software updates help reduce downtime and potential data loss when updating Pillar Axiom software and firmware modules. The system manages application dependencies using a compatibility matrix to ensure all required dependencies are met.

Administrators can schedule software updates up to 72 hours in advance, and the updates can be scheduled to occur during off-peak hours.

System Alerts

System alerts are messages that notify the administrator when a condition needs attention.

Many system alerts are accompanied by one or more menu selections that help the administrator resolve the condition.

Other system alerts clear automatically when the issue that generated the system alert has been resolved.

Pillar Axiom SMIProvider

Storage Management Initiative Specification (SMI-S) is a set of management interfaces designed to make it easier for storage hardware and management applications from different vendors to work together. Integral to this initiative, an SMI provider serves as a translator between the storage hardware and management applications. With the Pillar Axiom SMIProvider, any management application written natively to the SMI-S standard can manage Pillar Axiom systems.

SMIProvider can respond to requests from SMI-based management applications in supporting network and storage administrators who need to:

- Detect newly added devices on a network.
- Use encrypted SMI connections.
- Provision logical volumes and storage resources.
- Map HBAs to specific LUNs, mask LUNs from specific HBAs, or both.

The Pillar Axiom SMIProvider has the following features:

- Caches much of the information requested by SMI clients to provide enhanced response times.
- Supports Service Location Protocol device discovery using a built-in service agent.
- Supports NAS and SAN environments.
- Supports Fibre Channel protocols.
- Supports SMI version 1.1 (CTP 1.1 certified by SNIA), version 1.2, and version 1.3.

  Note: To access the services of the Pillar Axiom SMIProvider, the client needs to be compliant with SMI-S version 1.1 or later.

Pillar Axiom System Statistics

The Pillar Axiom system collects an extensive set of statistics that can be displayed in a variety of formats.

The complete set of statistics generated by the Pillar Axiom system can be parsed and formatted for import into the statistical tools of your choice. A subset of these statistics is available in statistical tables and pre-defined system reports in the graphical user interface (GUI) and command line interface (CLI). Many of these statistics can be displayed in trending charts in the GUI for a dynamic view of system interaction and performance.

Pillar Axiom Statistics Tools

The Pillar Axiom system collects statistics data on many facets of system operation. Each component of the Pillar Axiom system periodically collects key statistics covering such areas as filesystem read and write performance, block-level read and write performance, and error counts.

The Statistics Tools utilities make it possible to process and download any of these statistics from the Pillar Axiom graphical user interface (GUI) for analysis in statistical applications and spreadsheets.

Generated Reports

Generated reports provide listings of configuration details and statistical information about your system that you can download to your client from the Pillar Axiom system in various formats.

You can generate a statistical report immediately at the Generated Reports page, or you can schedule a report to be generated at a specified time at the Scheduled Reports page.

You can generate the following types of reports:

SAN Hosts	Provides statistical information on the host servers and configured components currently included in your storage area network (SAN).
Storage Performance	Provides performance information about the LUNs on the Pillar Axiom system. Includes operations/second, read MB/s, and write MB/s.

Storage Use	Provides storage capacity information on the storage currently available on the Pillar Axiom system. Includes total capacity, allocated, free, and unavailable capacity, and storage use by Storage Class.
Storage Use per Volume	Provides capacity information for each logical volume on the Pillar Axiom system.
System Configuration	Provides detailed information on the configuration and status of the current Pillar Axiom system and all of its components, such as serial numbers, firmware versions, ports, and status, for the Pilot, Slammers, and Bricks.
System Configuration Summary	Provides a summary of the Pilot, Slammer, and Brick information included in the detailed System Configuration report.

You can download reports in the following formats:

Table 5 Report download formats

Format	Usage
CSV	Comma-separated values. Import into any spreadsheet or database.
Excel	Import directly into an Excel spreadsheet.
HTML	Display in a browser.
PDF	Printer-friendly online document.
XML	XML tagged document.

Statistics and Trending Charts

The Pillar Axiom system collects statistics that can keep you informed about storage network status and performance.

Statistics and trending provides information about your storage area network (SAN) generated from the following:

- **LUN statistics**: Provides statistical information about LUNs such as capacity, throughput, I/O, and request time. Some of these statistics can be viewed in trending charts.

- **Slammer protocol statistics**: Provides Slammer port information such as speed and throughput, and the handling of commands and errors. Some of these statistics can be viewed in trending charts.

LUN statistics provide an overview of all LUNs present on your Pillar Axiom system. You can use this information to monitor the LUNs in your Pillar Axiom system or to compare LUN performance and capacity for planning improvements to your system.

You can create charts of statistics collected from the LUNs in your Pillar Axiom system to show trends in the statistics. Statistics collected from LUNs that can be used to produce trending charts include:

- The number and rates of data access commands handled over non-optimized paths.
- The number and rates of data access commands handled over optimized paths.
- Allocated capacity.

SAN Slammer protocol statistics provide an overview of the Slammer ports on your Pillar Axiom system.

You can create charts of statistics collected from the Slammer ports in your Pillar Axiom system to show trends in the data. Statistics collected from Slammer ports that can be used to produce trending charts include:

- The total amounts and rates of data read and written through a specified Slammer port.
- Statistics specific to the Fibre Channel (FC) or iSCSI protocol.

CHAPTER 4

NAS Overview

NAS Functionality

Pillar Axiom systems provide support for Network Attached Storage (NAS) Network File System (NFS) and Common Internet File System (CIFS) clients. NAS software runs on a NAS Slammer, which can service both NFS and CIFS requests.

Pillar Axiom systems feature a multi-node scalable NAS File Server with integrated high availability. Unlike many other NAS systems, the Pillar Axiom NAS File Server is able to achieve the highest possible throughput, fault tolerance, and uptime under many different failure scenarios. At the same time, it remains extremely easy to setup, configure, and maintain. By taking into consideration all the configuration options available for your network, you can achieve the highest possible throughput and redundancy.

Up to 1024 filesystems can be configured on the Pillar Axiom system at the same time.

In addition, Pillar Axiom systems allow NAS and SAN systems to share the same storage pool.

Network File System (NFS)

NFS is used primarily in UNIX environments to share filesystems across IP networks. Key features of the Pillar Axiom system's NFS implementation include:

- NFS version 2 and 3.
- Transmission Control Protocol (TCP) and User Datagram Protocol (UDP).
- Network Lock Manager (NLM) for advisory file-level and byte-range locking.
- Quotas for tree, account, and groups of accounts.

Common Internet File Systems (CIFS)

CIFS is used primarily in Microsoft Windows environments across IP networks. The Pillar Axiom system currently supports Microsoft Active Directory Services (ADS) using the Kerberos or NT LAN manager (NTLM) login sequence.

Key features of the CIFS implementation include:

- Share-level and byte-range locking.
- Opportunistic Locks—exclusive, batch, and level II.
- User authentication through consultation with a Windows Domain Controller or Active Directory Server, using Kerberos or NTLM protocol.
- Quotas for trees, accounts, and groups of accounts.

Concurrent NFS and CIFS Access

The Pillar Axiom system supports concurrent access to files through NFS and CIFS by coordinating locking and file permissions and by properly managing the differences between the NFS and CIFS protocols.

During concurrent access, Pillar Axiom system supports collaboration between NFS clients and CIFS clients by using the following rules:

- NFS write requests do respect CIFS mandatory locks.
- NFS read requests do not respect CIFS mandatory locks.
- NFS read and write requests adhere to CIFS ACL settings.
- CIFS read and write requests adhere to NFS file permissions.

Refer to the *Pillar Axiom CIFS and NFS Multi-Protocol Planning Guide* for details.

NAS Networking

The network layer of NAS implementation supports the following:

- Link aggregation (trunking). Pillar Axiom systems support the IEEE 802.3ad standard.

- Virtual local area networks (VLANs).
- Jumbo frames (MTU 9000 bytes).

Pillar Axiom SecureWORMfs

The Pillar Axiom SecureWORMfs is a type of filesystem used to enforce data retention. Data is stored on a SecureWORMfs in a non-erasable, non-rewritable (protected) manner. SecureWORMfs utilizes Write-Once-Read-Many (WORM) technology that permits use of a read-writable device to store data that, once set to protected, can never be modified and can only be erased when the retention period expires.

There are two types of SecureWORMfs filesystems:

- Standard: Data is retained on the filesystem for a fixed period of time that is specified by the retention period settings; however, data can be deleted at anytime by deleting the entire filesystem.
- Compliance: Stores critical business data as stipulated by various government regulations. Data is retained on the filesystem for a fixed period of time that is specified by file-level retention settings. A SecureWORMfs compliance filesystem cannot be deleted if there are protected files on the filesystem. To prevent malicious manipulation of the Pillar Axiom system clock, this feature requires that the Pillar Axiom system clock rely only on a customer-provided Network Time Protocol (NTP) server.

SecureWORMfs retention is implemented when a protected file is closed.

Note: Pillar Axiom release 4.0 and later permits customers to downgrade a SecureWORMfs compliance filesystem to a SecureWORMfs standard filesystem so that it can be deleted. To delete the SecureWORMfs, the customer must contact the Pillar World Wide Customer Support Center. For your protection, authentication will be required.

CHAPTER 5

SAN Overview

SAN Functionality

LUNs are accessed by host systems through a storage area network (SAN) attachment. The Pillar Axiom system supports Fibre Channel (FC) host bus adapter (HBA) and iSCSI HBA SAN attachments. Supported host operating systems include the following:

- AIX
- HP-UX
- Linux
- Solaris
- Windows

See the *Pillar Axiom Interoperability Guide* for a current list of supported host operating systems.

Up to 8192 initiators (host systems) of any of the supported operating systems can be connected to the Pillar Axiom system at the same time.

Supported FC topologies include the following:

- Fabric
- Private FC-AL
- Point-to-point
- Public loop

Supported iSCSI features include:

- Optional Challenge-Handshake Authentication Protocol (CHAP) authentication of iSCSI initiators.
- Configurable Pillar Axiom CHAP secret for bi-directional CHAP support.
- Optional access control of iSCSI initiators.
- Explicit or DHCP-based IP address configuration.

- Internet Storage Name Service (iSNS).
- Use of a RADIUS server for iSCSI authentication.

Refer to the *Pillar Axiom iSCSI Integration Guide for Windows Platforms* for details.

The Pillar Axiom system supports a maximum of:

- 16384 LUNs for each system.
- 256 LUNs for each host.
- 8192 LUNs for each SAN Slammer.
- 256 TCP connections for each iSCSI port.
- 256 iSCSI Initiators for each iSCSI port.
- 32 persistent reservation registration keys for each LUN.
- 512 simultaneous commands for each iSCSI port.
- 2048 simultaneous commands for each FC port.

Pillar Axiom Path Manager

A host system communicates with LUNs through a normal SCSI-over-Fibre Channel (FC) or Internet Small Computer System Interface (iSCSI) initiator driver. Typically, these drivers are provided with the operating system or the HBA manufacturer.

For many operating systems, we supply an optional component called the Pillar Axiom Path Manager (APM) to run on the host. APM performs many key functions:

- Enables the use of multiple paths between HBA ports on the host and Slammer ports.
- Balances the load across HBA channels and Slammer ports.
- Monitors path failure and transfers traffic to different paths if a failure occurs.
- Maps LUNs to host volumes.
- Automates the recognition and configuration of the host in the Pillar Axiom Storage Services Manager.
- Reports the status of host drivers and paths in the Pillar Axiom Storage Services Manager.
- Collects logs to include in Call-Home files when SAN hosts are selected.

LUN Configuration

SAN Slammers support both LUN mapping and LUN masking. LUN mapping presents a LUN to a restricted set of hosts as any LUN number. LUN masking can restrict host access to only a defined set of Slammer ports. Initiator access to LUNs is controlled through the initiator's World Wide Port Name (WWPN) or Internet Small Computer System Interface (iSCSI) initiator name. A maximum of 256 LUNs may be exposed to a single host.

Individual LUNs can have different Quality of Service (QoS) settings (such as performance, redundancy, I/O bias, and striping). Multiple LUNs can be mapped to multiple different hosts and multiple hosts can be mapped to an individual LUN, regardless of their respective QoS settings.

iSNS

The Internet Storage Name Service (iSNS) facilitates automated discovery, management, and configuration of iSCSI devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in Fibre Channel networks, allowing a commodity IP network to function in a capacity similar to that of a storage area network.

The iSNS feature expects all Pillar Axiom iSCSI ports to have access to the same primary iSNS server. This rule is necessary so that all iSCSI ports can expect the same result when querying the iSNS database for the set of initiators that are members of the Pillar Axiom Discovery Domain Set.

CHAPTER 6

Remote Replication Overview

About Remote Replication

The Pillar Axiom system provides replication for both network attached storage (NAS) and storage area network (SAN) environments. The remote replication feature provides efficient operational or disaster recovery (DR) capabilities that improve business data recovery while allowing you to adhere to regulatory or compliance requirements without complex configuration.

Pillar Axiom MaxRep replication supports short recovery time objectives (RTOs) by making the replicated volumes immediately available, allowing client applications to quickly switch over to and begin accessing the replicated data. When the disaster has passed and the original Pillar Axiom system has been repaired, you can efficiently and quickly restore the original volumes on the repaired system.

Pillar Axiom MaxRep replication uses block-level operations to transfer data during replication operations over local and wide area networks. Pillar Axiom systems help guard against security exposures by requiring passwords for all replication sessions. If you need security on both ends of the communication link between two Pillar Axiom systems engaged in data transfer operations, you will need to install the appropriate networking hardware to encrypt the data and otherwise secure that link.

Pillar Axiom MaxRep Replication for NAS provides data protection for filesystems. A source filesystem can be replicated to multiple different targets, if desired. Depending on your recovery point objectives (RPOs), synchronizing the source filesystem with any or all of the target filesystems can be performed quickly.

Pillar Axiom MaxRep Replication for SAN provides complete data protection and disaster recovery for SAN LUNs. A source LUN can be replicated to the same Pillar Axiom system or to a different Pillar Axiom system (located within the local Metro area or at a remote site). The destination LUN is updated synchronously or asynchronously, depending on the configuration. Optionally, for supported applications checkpoints can be taken to ensure consistency of application data on the source and destination LUNs at any point in time. In cases of data corruption, it is possible for the user to roll back to a point in time prior to when the corruption took place.

Pillar Axiom MaxRep for NAS

Pillar Axiom MaxRep Replication for NAS is a native command-line utility that provides data replication among Pillar Axiom systems in a network attached storage (NAS) environment.

Typical usage of this software might transpire in this way:

- Establish a relationship between two compatible filesystems. In this relationship, one filesystem is the source of replication and the other filesystem is the target.
- Fully synchronize the content of the two filesystems.
- Perform an incremental synchronization some number of times.
- Break the relationship.
- Mark the target filesystem as being live (readable and writable).
- Use the target filesystem in the same way as any other filesystem.

Note: All replication configuration and management is performed on the command line.

Using block-based read and write operations to perform the replication, the Pillar Axiom MaxRep Replication for NAS utility simulates a backup operation on the source filesystem and a restore operation on the target filesystem. (See the following figure.) The source and target filesystems can reside on the same Pillar Axiom system or on two separate systems that are connected over a local or wide area network.

Replication periodically saves specific information that allows replication to restart as needed. This information is called a *restart point*. For example, if, while the Pillar Axiom software is being updated, the system warmstarts or restarts during the software update, replication uses the restart point to continue the replication operations after the system returns to a Normal state.

Figure 10 System hardware in a NAS replication environment

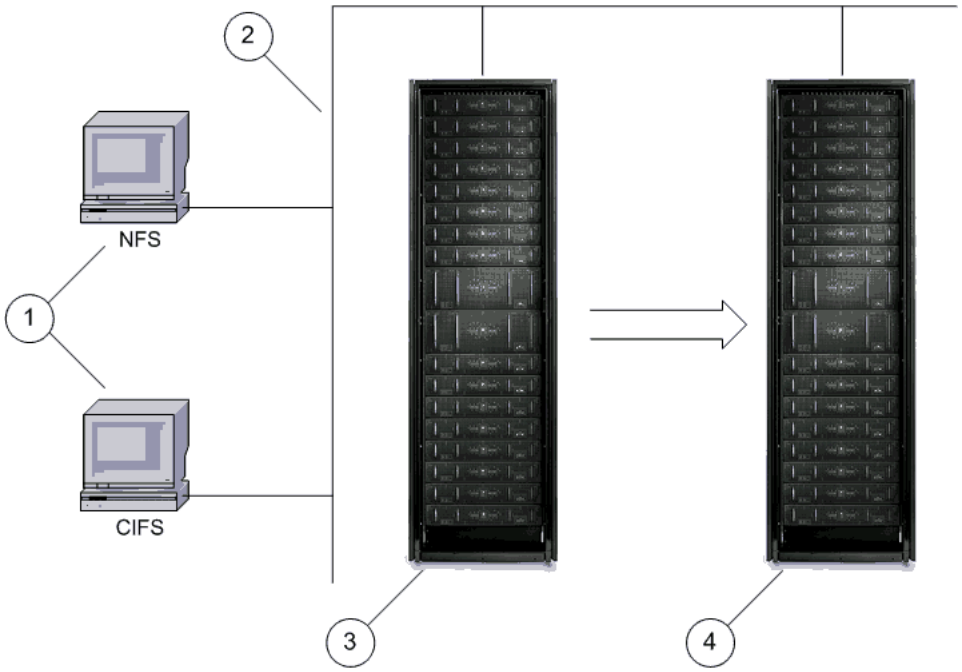

Legend		
	1 Client hosts	3 Pillar Axiom system A (source)
	2 TCP/IP network	4 Pillar Axiom system B (target)

The Pillar Axiom MaxRep Replication for NAS utility includes client software for host machines that use either Common Internet File System (CIFS) or Network File System (NFS) protocols to communicate with Pillar Axiom servers. You can use the client software for initiating replication operations and for otherwise managing replication objects. This utility allows you to:

- Discover information that is essential to replication operations (such as determining the filesystems that are available for replication).
- Establish a relationship between two filesystems.
- Synchronize two filesystems that have been defined as a replication pair. The synchronization is a controlled, highly organized transfer of data from the source filesystem to the target.
- Place the target filesystem into a *live* mode to make it available for read and write operations by users.

  If the source filesystem fails, the administrator can change the state of the target filesystem from its special passive mode to an active mode that supports user I/O.

To support replication, the system administrator must enable Network Data Management Protocol (NDMP) and configure an NDMP File Server on both the source and the target Pillar Axiom systems. Ideally, this File Server should be dedicated to replication and other NDMP operations and not used for regular Common Internet File System (CIFS) and Network File System (NFS) user traffic.

For performance reasons, an NDMP File Server should have multiple virtual interfaces (VIFs) associated with it, one for each filesystem participating in replication. (A VIF configuration defines a physical network port, including an IP address, on a Slammer control unit (CU). During replication, data transfer occurs through these ports.)

Because a filesystem is homed on a particular Slammer CU, the administrator should consider configuring one NDMP VIF for each NAS Slammer CU that is home to a filesystem that participates in replication. If an NDMP VIF is not defined, the replication software uses an existing VIF that is associated with the filesystem, possibly competing with regular NFS and CIFS user traffic.

NAS replication uses only those VIFs that have been defined for the NDMP File Server. Furthermore, the utility automatically selects the VIFs that are most optimal, which could result in a different VIF for each replication connection.

Note: NDMP traffic and replication traffic share the same File Server and VIFs.

For information on configuring NDMP, File Servers, and VIFs in a Pillar Axiom system, refer to the *Pillar Axiom Administrator's Guide*.

Pillar Axiom MaxRep for SAN

Pillar Axiom MaxRep Replication for SAN enables you to automatically replicate and restore Pillar Axiom system data in a storage area network (SAN) environment.

In SAN replication, one or more pairs of parallel volumes, called *replication pairs*, are established at primary and secondary locations, typically on separate remotely distributed Pillar Axiom systems. Communication links are established between these primary and secondary locations to maintain consistency of the data on the volumes at both locations. The transfer of data takes place automatically as the data on the source volume changes at the primary site. Those changes are replicated to the destination volume at the secondary site. The replication pair updates continuously as long as the integrity of both volumes persists and the communication link between the locations is maintained.

A *replication pair* is the combination of a source and target LUN that are associated with each other. These LUNs can reside on the same Pillar Axiom system, or located on separate Pillar Axiom systems at the same geographical location, or on separate, remotely distributed Pillar Axiom systems. The

replication process requires one or more Pillar Axiom Replication Engines. Communication links are established between the Pillar Axiom Replication Engines and the primary and secondary locations to provide a reliable transport mechanism. Data transfer takes place automatically when the data stored on the primary location changes, and these changes are duplicated on the volume at the secondary location.

Pillar Axiom MaxRep Replication for SAN supports synchronous or asynchronous replication of LUNs or application consistent volume sets. When the primary and secondary locations are geographically distributed and communication is over a wide area network (WAN) link, separate primary and secondary Pillar Axiom Replication Engines at both the locations are required. To ensure high availability, two Pillar Axiom Replication Engines can be clustered at one or both locations.

Figure 11 Asynchronous Pillar Axiom MaxRep Replication for SAN configuration

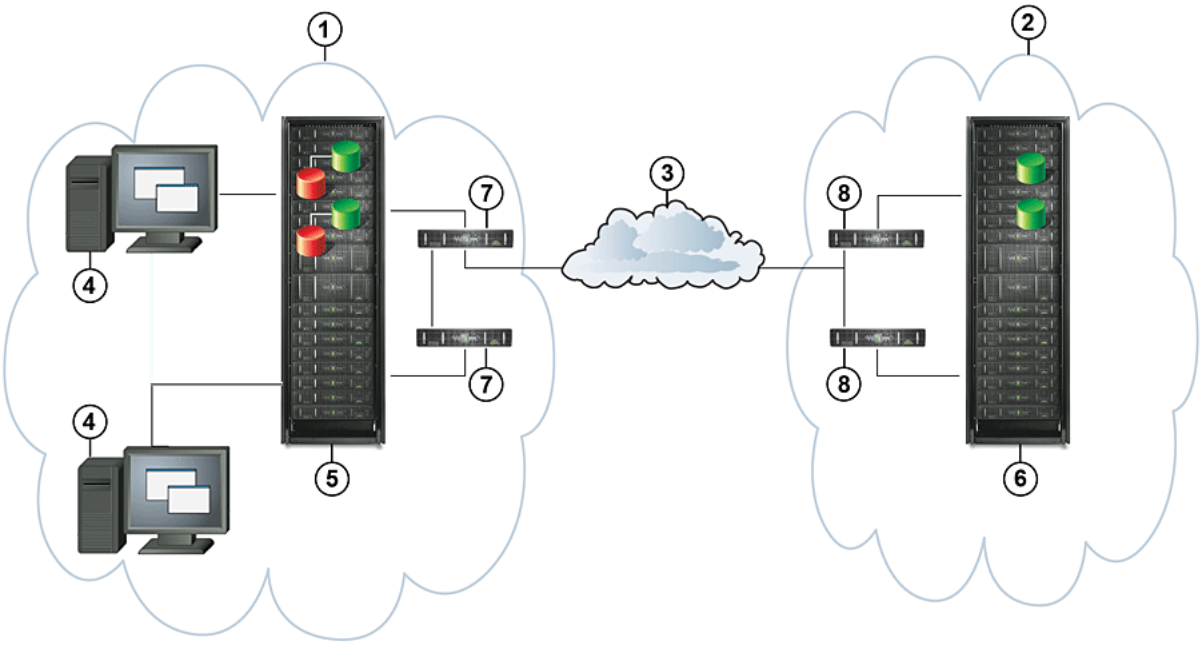

Legend		
	1 Primary site	5 Primary Pillar Axiom system
	2 Secondary site	6 Secondary Pillar Axiom system
	3 WAN connection	7 Pillar Axiom Replication Engines on the primary site clustered for high availability
	4 Application servers	8 Pillar Axiom Replication Engines on the secondary site clustered for high availability

Data can be recovered from either the primary or the secondary site, and the direction of replication can be reversed. Several failover and failback scenarios

can be planned and implemented using Pillar Axiom MaxRep Replication for SAN.

CHAPTER 7

Local Data Protection Overview

Data Protection Services

A Pillar Axiom system provides the following forms of replication, each with its own distinct purpose:

- Snap FS—a read-only filesystem-based snapshot.
- Clone FS—a read-write filesystem-based snapshot.
- Clone LUN—a read-write LUN-based snapshot.
- Pillar Axiom Data Protection Manager snapshot—a snapshot of data from a Microsoft Volume Shadow Copy Service (VSS)-enabled application.
- Volume Copy—a read-write full volume replica.

Partial-image snapshots (Clone FS and Clone LUN) require storage space allocated for clones. This space is allocated at the time the volume is created. Only changes made to the source volume or the clone are stored on the system in the storage space allocated for clones.

Partial-image snapshots are recommended for short-term backups when you expect only a moderate amount of data to change.

Full-image snapshots (Volume Copies) are the same size as the source volume. The system stores a complete copy of the source volume from the time the snapshot was taken. The data is copied from the source to the snapshot in the background and the full-image snapshot automatically becomes a regular standalone volume once all of the data has been copied.

Full-image snapshots are recommended when a large amount of data will be changing or if you know that the snapshot will be needed as a standalone volume.

Snap FS

Snap FS is a filesystem snapshot that serves primarily as a file recovery mechanism. Snapshots are point-in-time copies that use the active filesystem as a base. A Snap FS preserves a view of the data at the exact time of the

snapshot, allowing users to access older versions of files from a hidden subdirectory within the primary filesystem. Snap FSs are read-only.

A common reason for making periodic Snap FSs is to allow ordinary users to recover files that have been accidentally deleted without any assistance from the Pillar Axiom system administrator.

Another snapshot scenario is when a system administrator changes a software program and then finds the changes do not work. A snapshot allows the administrator to return to a previous version of the program and protects against inadvertent deletion of files.

The Pillar Axiom Storage Services Manager (the GUI), allows you to build and manage snapshot (Snap FS) schedules. A typical approach is to perform snapshots hourly and retain those snapshots for a day or a week.

Clone FS

Clone FS is a point-in-time, read-write copy of a filesystem that you intend to snap (split) from the source filesystem for immediate use. You can create an immediate snapshot of a filesystem. If a filesystem requires a filesystem check (FSCK), you can clone the filesystem, FSCK the clone, and then synchronize the two.

Clone FSs utilize partial-block snapshot technology. Because Clone FSs point to the original data, and have the same QoS parameters as the source filesystem, you cannot change the Quality of Service (QoS) attributes or the Clone FS location on the Slammer control unit (CU). Only data that has been modified on either the source filesystem or the Clone FS is copied. This means that Clone FSs can use significantly smaller storage space than full volume copies.

Clone LUN

Clone LUN is a point-in-time copy of a LUN that you intend to snap (split) from the source LUN for immediate read-write use. You can create an immediate snapshot of a LUN. Clone LUNs utilize partial-block snapshot technology.

Clone LUNs point to the original data and are created using the same QoS parameters as the source LUN. You can, however, change the priority level Quality of Service (QoS) attribute after you create the Clone LUN. Only data that has been modified on either the source LUN or the Clone LUN is stored in the clone storage space. This means that Clone LUNs can use significantly smaller storage space than full volume copies.

Volume Copy

A full-volume copy is a point-in-time, block-for-block replication of the source volume (filesystem or LUN). You can use the copy for backup, testing, reporting, or data warehousing. These workloads are directed to the copied volume, not to the source. Thus, the primary volume can service application I/Os with minimal performance degradation from ancillary activities. Volume Copies have read and write capabilities.

The Volume Copy is made to a new volume which can have its own QoS metrics. This allows system resources to be maximized for the task at hand. For example, a replicated volume that is used for reporting is assigned a lower performance priority and a higher read-centric access pattern than would the source volume.

Clone Storage Space

Clones consume space on the system. This is referred to as the allocated clone storage space or clone repository.

The system allocates space for clones during volume creation. Sufficient space for the clones must be available on the system. The amount of allocated clone space will depend on the expected amount of change to the volume and the frequency at which clones will be created. The system stores only the changes made to either the source volume or the clone in the allocated clone storage space.

Data Protection Manager

The Pillar Axiom Data Protection Manager (DPM) is an application that manages application-aware backup and recovery for the Pillar Axiom 600 system. The Data Protection Manager 2.0 runs on Windows Server 2003 or 2008, which may be run from a physical or virtual environment. To work in a virtual environment, DPM requires Pillar Axiom system release 5.0, or later.

In Microsoft environments, Data Protection Manager maintains data integrity by using the Microsoft Volume Shadow Copy Service (VSS) provider. The VSS infrastructure enables data protection and recovery without interrupting normal operations.

DPM gives you the capability to schedule backups, called checkpoints, of the application data on a regular basis. You can control the number of checkpoints that the system retains by using the DPM retention policy. The retention policy allows you to maintain the maximum number of checkpoints or set a maximum amount of time to keep older checkpoints on your system. You can override this policy on selected checkpoints.

The Data Protection Manager uses a Microsoft VSS technique called LUN Resync to restore the source LUNs. The restore process uses the Clone LUNs on the Pillar Axiom system to restore the LUNs.

Reduced Clone LUN Overhead

Most snapshot implementations handle the necessary data manipulation as part of the write process. Pillar Axiom SAN systems handle the cloned data processing during the write de-stage process as data is flushed from protected storage to the physical drives. This technique enables a write request to complete more quickly. This implementation allows for the use of Clone LUNs with minimal degradation of system performance.

Pillar Axiom SnapDelta FS

The Pillar Axiom SnapDelta FS feature lists files that have been created, renamed, or deleted, or files with changed content, within a Pillar Axiom filesystem during the interval between two Snap FS snapshots.

External applications sometimes need to scan the contents of a Pillar Axiom filesystem to extract information. After an initial scan, these applications need to perform periodic rescans to process new, renamed, and deleted files, and files with changed content. Examples of these external applications include:

- File-based replication applications.
- Search and indexing applications.
- Information classification applications.
- Virus scanning applications.

The Pillar Axiom 600 storage system provides a filesystem change reporting command, `axiom_snapdelta_fs`, to enable external applications to rescan Pillar Axiom filesystems. SnapDelta FS provides efficient access to the set of file changes created during the interval between two Snap FS snapshots.

For example, here is how an external application might use SnapDelta FS to process filesystem changes:

- First, an application that tracks changes to a filesystem creates a snapshot and performs a full scan of the contents of the filesystem in that snapshot.
- Later, the application creates a second snapshot and uses the `axiom_snapdelta_fs` command to find the changes between the first snapshot and the second snapshot.
- Next, the application creates a third snapshot and performs further tracking of filesystem changes by finding the changes between the second and third snapshot, and so on.

Filesystem change reporting works on two snapshots at a time. These snapshots may be immediate snapshots created specifically for the use of `axiom_snapdelta_fs`, or they may be scheduled snapshots created automatically for the filesystem.

Backup Tools

The Pillar Axiom system provides tools to facilitate backup. The following sections describe these tools.

NDMP-Based Backup System

Network Data Management Protocol (NDMP) is an industry-standard protocol that allows for the use of third-party backup applications to manage the backup and recovery of customer data. An NDMP user account, password, and access port are configured through the Pilot. Pillar Axiom systems support NDMP version 4. Refer to http://www.ndmp.org/info/faq.shtml for details.

NDMP-based backup and restore operations can be integrated into your existing backup and recovery system. When you do this, you can completely automate the backup and restore operations.

The Pillar Axiom system supports:

- CommVault Galaxy
- Symantec Veritas NetBackup
- EMC NetWorker
- BakBone NetVault: BackUp
- Oracle Secure Backup

Note: Refer to the *Pillar Axiom Support and Interoperability Guide* for more details on the supported versions of the above mentioned backup and recovery software.

Microsoft Volume Shadow Copy Service (VSS)

Pillar Axiom systems can use the Microsoft Volume Shadow Copy Service (VSS) in their backup and restore solutions.

VSS enables data protection and management services through a standard set of configuration and monitoring capabilities. These capabilities include creating, manipulating, and restoring snapshots without shutting down applications or essential services.

Index